

On Hölder projective divergences*

Frank Nielsen[†]

École Polytechnique, France
Sony Computer Science Laboratories, Japan

Ke Sun

Computer, Electrical and Mathematical Sciences and Engineering Division
King Abdullah University of Science and Technology (KAUST), Saudi Arabia

Stéphane Marchand-Maillet

Viper Group, Computer Vision and Multimedia Laboratory
University of Geneva, Switzerland

Abstract

We describe a framework to build distances by measuring the tightness of inequalities, and introduce the notion of proper statistical divergences and improper pseudo-divergences. We then consider the Hölder ordinary and reverse inequalities, and present two novel classes of Hölder divergences and pseudo-divergences that both encapsulate the special case of the Cauchy-Schwarz divergence. We report closed-form formulas for those statistical dissimilarities when considering distributions belonging to the same exponential family provided that the natural parameter space is a cone (e.g., multivariate Gaussians), or affine (e.g., categorical distributions). Those new classes of Hölder distances are invariant to rescaling, and thus do not require distributions to be normalized. Finally, we show how to compute statistical Hölder centroids with respect to those divergences, and carry out center-based clustering toy experiments on a set of Gaussian distributions that demonstrate empirically that symmetrized Hölder divergences outperform the symmetric Cauchy-Schwarz divergence.

Keywords: *Hölder inequalities; Hölder divergences; projective divergences; Cauchy-Schwarz divergence; Hölder escort divergences; skew Bhattacharyya divergences; exponential families; conic exponential families; escort distribution; clustering.*

1 Introduction: Inequality, proper divergence and improper pseudo-divergence

1.1 Statistical divergences from inequality gaps

An inequality is denoted mathematically by $\text{lhs} \leq \text{rhs}$, where lhs and rhs denote respectively the *left-hand-side* and *right-hand-side* of the inequality. One can build *dissimilarity* measures from inequalities $\text{lhs} \leq \text{rhs}$ by measuring the *inequality tightness*: For example, we may quantify the tightness of an inequality by its *difference gap*:

$$\Delta = \text{rhs} - \text{lhs} \geq 0. \quad (1)$$

When $\text{lhs} > 0$, the inequality tightness can also be gauged by the *log-ratio gap*:

$$D = \log \frac{\text{rhs}}{\text{lhs}} = -\log \frac{\text{lhs}}{\text{rhs}} \geq 0. \quad (2)$$

We may further compose this inequality tightness value measuring non-negative gaps with a *strictly monotonically increasing function* f (with $f(0) = 0$).

*Reproducible source code available at <https://www.lix.polytechnique.fr/~nielsen/HPD/>

[†]Contact author: Frank.Nielsen@acm.org

A *bi-parametric* inequality $\text{lhs}(p, q) \leq \text{rhs}(p, q)$ is called *proper* if it is strict for $p \neq q$ (i.e., $\text{lhs}(p, q) < \text{rhs}(p, q), \forall p \neq q$) and *tight* if and only if (iff) $p = q$ (i.e., $\text{lhs}(p, q) = \text{rhs}(p, q), \forall p = q$). Thus a proper bi-parametric inequality allows one to define dissimilarities such that $D(p, q) = 0$ iff $p = q$. Such a dissimilarity is called *proper*. Otherwise, an inequality or dissimilarity is said *improper*. Note that there are many equivalent words used in the literature instead of (dis-)similarity: distance (although often assumed to have metric properties), pseudo-distance, discrimination, proximity, information deviation, etc.

A *statistical dissimilarity* between two discrete or continuous distributions $p(x)$ and $q(x)$ on a support \mathcal{X} can thus be defined from inequalities by summing up or taking the integral for the inequalities instantiated on the observation space \mathcal{X} :

$$\forall x \in \mathcal{X}, \quad D_x(p, q) = \text{rhs}(p(x), q(x)) - \text{lhs}(p(x), q(x)) \Rightarrow \quad (3)$$

$$D(p, q) = \begin{cases} \sum_{x \in \mathcal{X}} [\text{rhs}(p(x), q(x)) - \text{lhs}(p(x), q(x))] & \text{discrete case,} \\ \int_{\mathcal{X}} [\text{rhs}(p(x), q(x)) - \text{lhs}(p(x), q(x))] dx & \text{continuous case.} \end{cases} \quad (4)$$

In such a case, we get a *separable divergence*. Some non-separable inequalities induce a *non-separable divergence*. For example, the renown Cauchy-Schwarz divergence [8] is not separable because in the inequality:

$$\int_{\mathcal{X}} p(x)q(x)dx \leq \sqrt{\left(\int_{\mathcal{X}} p(x)^2 dx\right) \left(\int_{\mathcal{X}} q(x)^2 dx\right)}, \quad (5)$$

the rhs is not separable.

Furthermore, a proper dissimilarity is called a *divergence* in information geometry [2] when it is C^3 (i.e., three times differentiable thus allowing to define a metric tensor [34] and a cubic tensor [2]).

Many familiar distances can be reinterpreted as inequality gaps in disguise. For example, Bregman divergences [4] and Jensen divergences [28] (also called Burbea-Rao divergences [9, 22]) can be reinterpreted as inequality difference gaps, and the Cauchy-Schwarz distance [8] as an inequality log-ratio gap:

Example 1 (Bregman divergence as a Bregman score-induced gap divergence). *A proper score function [14] $S(p : q)$ induces a gap divergence $D(p : q) = S(p : q) - S(p : p) \geq 0$. A Bregman divergence [4] $B_F(p : q)$ for a strictly convex and differentiable real-valued generator $F(x)$ is induced by the Bregman score $S_F(p : q)$. Let $S_F(p : q) = -F(q) - \langle p - q, \nabla F(q) \rangle$ denote the Bregman proper score minimized for $p = q$. Then the Bregman divergence is a gap divergence: $B_F(p : q) = S_F(p : q) - S_F(p : p) \geq 0$. When F is strictly convex, the Bregman score is proper, and the Bregman divergence is proper.*

Example 2 (Cauchy-Schwarz distance as a log-ratio gap divergence). *Consider the Cauchy-Schwarz inequality $\int_{\mathcal{X}} p(x)q(x)dx \leq \sqrt{\left(\int_{\mathcal{X}} p(x)^2 dx\right) \left(\int_{\mathcal{X}} q(x)^2 dx\right)}$. Then the Cauchy-Schwarz distance [8] between two continuous distributions is defined by $\text{CS}(p(x) : q(x)) = -\log \frac{\int_{\mathcal{X}} p(x)q(x)dx}{\sqrt{\left(\int_{\mathcal{X}} p(x)^2 dx\right) \left(\int_{\mathcal{X}} q(x)^2 dx\right)}} \geq 0$.*

Note that we use the modern notation $D(p(x) : q(x))$ to emphasize that the divergence is potentially asymmetric: $D(p(x) : q(x)) \neq D(q(x) : p(x))$, see [2]. In information theory [11], the older notation “||” is often used instead of “:” that is used in information geometry [2].

To conclude this introduction, let us finally introduce the notion of *projective statistical distances*. A statistical distance $D(p(x) : q(x))$ is said *projective* when:

$$D(\lambda p(x) : \lambda' q(x)) = D(p(x) : q(x)), \quad \forall \lambda, \lambda' > 0. \quad (6)$$

The Cauchy-Schwarz distance is a projective divergence. Another example of such a projective divergence is the parametric γ -divergence [13].

Example 3 (γ -divergence as a projective score-induced gap divergence). *The γ -divergence [13, 29] $D_{\gamma}(p(x) : q(x))$ for $\gamma > 0$ is projective:*

$$D_{\gamma}(p(x), q(x)) = S_{\gamma}(p(x), q(x)) - S_{\gamma}(p(x), p(x)), \text{ with} \\ S_{\gamma}(p(x), q(x)) = -\frac{1}{\gamma(1+\gamma)} \frac{\int p(x)q(x)^{\gamma} dx}{\left(\int q(x)^{1+\gamma} dx\right)^{\frac{\gamma}{1+\gamma}}}.$$

The γ -divergence is related to the proper pseudo-spherical score [13].

The γ -divergences have been proven useful for robust statistical inference [13] in the presence of heavy outlier contamination.

1.2 Pseudo-divergences and the axiom of indiscernability

Consider a broader class of *statistical pseudo-divergences* based on *improper inequalities*, where the tightness of $\text{lhs}(p, q) \leq \text{rhs}(p, q)$ does not imply that $p = q$. This family of dissimilarity measures have interesting properties which have not been studied before.

Formally, statistical pseudo-divergences are defined with respect to density measures $p(x)$ and $q(x)$ with $x \in \mathcal{X}$, where \mathcal{X} denotes the support. By definition, pseudo-divergences satisfy the following three fundamental properties:

1. Non-negativeness: $D(p(x) : q(x)) \geq 0$ for any $p(x), q(x)$;
2. Reachable indiscernability:
 - $\forall p(x)$, there exists $q(x)$ such that $D(p(x) : q(x)) = 0$,
 - $\forall q(x)$, there exists $p(x)$ such that $D(p(x) : q(x)) = 0$.
3. Positive correlation: If $D(p : q) = 0$, then $(p(x_1) - p(x_2))(q(x_1) - q(x_2)) \geq 0$ for any $x_1, x_2 \in \mathcal{X}$.

As compared to *statistical divergence* measures such as the Kullback-Leibler (KL) divergence:

$$\text{KL}(p(x) : q(x)) = \int_{\mathcal{X}} p(x) \log \frac{p(x)}{q(x)} dx, \quad (7)$$

pseudo-divergences do not require $D(p(x) : p(x)) = 0$. Instead, any pair of distributions $p(x)$ and $q(x)$ with $D(p(x) : q(x)) = 0$ only has to be “positively correlated” such that $p(x_1) \leq p(x_2)$ implies $q(x_1) \leq q(x_2)$, and vice versa. Any divergence with $D(p(x) : q(x)) = 0 \Rightarrow p(x) = q(x)$ (law of indiscernibles) automatically satisfies this weaker condition, and therefore any divergence belongs to the broader class of pseudo-divergences. Indeed, if $p(x) = q(x)$ then $(p(x_1) - p(x_2))(q(x_1) - q(x_2)) = (p(x_1) - p(x_2))^2 \geq 0$. However the converse is not true. As we shall describe in the remainder, the family of pseudo-divergences is *not* limited to proper divergence measures. In the remainder, the term “pseudo-divergence” refers to such divergences that are *not* proper divergence measures.

We study two novel statistical dissimilarity families: One family of statistical improper pseudo-divergences and one family of proper statistical divergences. Within the class of pseudo-divergences, this work concentrates on defining a one-parameter family of dissimilarities called Hölder log-ratio gap divergence that we concisely abbreviate as HPD for “Hölder pseudo divergence” in the remainder. We also study its proper divergence counterpart termed HD for “Hölder divergence.”

1.3 Prior work and contributions

The term “Hölder divergence” has first been coined in 2014 based on the definition of the *Hölder score* [20, 19]: The score-induced Hölder divergence $D(p(x) : q(x))$ is a proper gap divergence that yields a scale-invariant divergence [20, 19]. Let $p_{a,\sigma}(x) = a\sigma p(\sigma x)$ for $a, \sigma > 0$ be a transformation. Then a scale-invariant divergence [20, 19] satisfies $D(p_{a,\sigma}(x) : q_{a,\sigma}(x)) = \kappa(a, \sigma)D(p(x) : q(x))$ for a function $\kappa(a, \sigma) > 0$. This gap divergence is proper since it is based on the so-called Hölder score [20, 19] but is *not* projective and does not include the Cauchy-Schwarz divergence. Due to these differences the Hölder log-ratio gap divergence introduced here shall not be confused with the Hölder gap divergence induced by the *Hölder score* [19, 20] that relies both on a scalar γ and a function $\phi(\cdot)$.

We shall introduce *two* novel families of log-ratio projective gap divergences based on Hölder ordinary (or forward) and reverse inequalities that extend the Cauchy-Schwarz divergence, study their properties, and consider as an application clustering Gaussian distributions: We experimentally show better clustering results when using symmetrized Hölder divergences than using the Cauchy-Schwarz divergence. To contrast with the “Hölder composite score-induced divergences” of [19], our Hölder divergences admit closed-form expressions between distributions belonging to the same exponential families [23] provided that the natural parameter space is a cone or affine.

Our main contributions are summarized as follows:

- Define the uni-parametric family of Hölder improper pseudo-divergences (HPDs) in §2 and the bi-parametric family of Hölder proper divergences in §3 (HDs) for positive and probability measures, and study their properties (including their relationships with skewed Bhattacharyya distances [22] via escort distributions);
- Report closed-form expressions of those divergences for exponential families when the natural parameter space is a cone or affine (include but not limited to the cases of categorical distributions and multivariate Gaussian distributions) in §4;
- Provide approximation techniques to compute those divergences between mixtures based on log-sum-exp inequalities in §4.6;
- Describe a variational center-based clustering technique based on the convex-concave procedure for computing Hölder centroids, and report our experimental results in §5.

1.4 Organization

This paper is organized as follows: §2 introduces the definition and properties of Hölder pseudo-divergences (HPDs). It is followed by §3 that describes Hölder proper divergences (HDs). In §4, closed-form expressions for those novel families of divergences are reported for the categorical, multivariate Gaussian, Bernoulli, Laplace and Wishart distributions. §5 defines Hölder statistical centroids and presents a variational k -means clustering technique: We show experimentally that using Hölder divergences improve over the Cauchy-Schwarz divergence. Finally, §6 concludes this work and hints at further perspectives from the viewpoint of statistical estimation and manifold learning. In Appendix A, we recall the proof of the ordinary and reverse Hölder’s inequalities.

2 Hölder pseudo-divergence: Definition and properties

Hölder’s inequality (see [18] and Appendix A for a proof) states for positive real-valued functions¹ $p(x)$ and $q(x)$ defined on the support \mathcal{X} that:

$$\int_{\mathcal{X}} p(x)q(x)dx \leq \left(\int_{\mathcal{X}} p(x)^{\alpha} dx \right)^{\frac{1}{\alpha}} \left(\int_{\mathcal{X}} q(x)^{\beta} dx \right)^{\frac{1}{\beta}}, \quad (8)$$

where exponents α and β satisfy $\alpha\beta > 0$ as well as the *exponent conjugacy* condition: $\frac{1}{\alpha} + \frac{1}{\beta} = 1$. We also write $\beta = \bar{\alpha} = \frac{\alpha}{\alpha-1}$ meaning that α and β are *conjugate* Hölder exponents. We check that $\alpha > 1$ and $\beta > 1$. Hölder inequality holds even if the lhs is infinite (meaning that the integral diverges) since the rhs is also infinite in that case.

The *reverse Hölder inequality* holds for conjugate exponents $\frac{1}{\alpha} + \frac{1}{\beta} = 1$ with $\alpha\beta < 0$ (then $0 < \alpha < 1$ and $\beta < 0$, or $\alpha < 0$ and $0 < \beta < 1$):

$$\int_{\mathcal{X}} p(x)q(x)dx \geq \left(\int_{\mathcal{X}} p(x)^{\alpha} dx \right)^{\frac{1}{\alpha}} \left(\int_{\mathcal{X}} q(x)^{\beta} dx \right)^{\frac{1}{\beta}}. \quad (9)$$

Both Hölder’s inequality and the reverse Hölder inequality turn tight when $p(x)^{\alpha} \propto q(x)^{\beta}$ (see proof in Appendix A).

2.1 Definition

Let $(\mathcal{X}, \mathcal{F}, \mu)$ be a measurable space where μ is the Lebesgue measure, and let $L^{\gamma}(\mathcal{X}, \mu)$ denote the Lebesgue space of functions that have their γ -th power of absolute value Lebesgue integrable, for any $\gamma > 0$ (when $\gamma \geq 1$, $L^{\gamma}(\mathcal{X}, \mu)$ is a Banach space). We define the following pseudo-divergence:

¹In a more general form, Hölder’s inequality holds for any real and complex valued functions. In this work, we only focus on real positive functions that are densities of positive measures.

Definition 1 (Hölder statistical pseudo-divergence, HPD). *For conjugate exponents α and β with $\alpha\beta > 0$, the Hölder Pseudo-Divergence (HPD) between two densities $p(x) \in L^\alpha(\mathcal{X}, \mu)$ and $q(x) \in L^\beta(\mathcal{X}, \mu)$ of positive measures absolutely continuous with respect to (wrt.) μ is defined by the following log-ratio gap:*

$$D_\alpha^H(p(x) : q(x)) = -\log \left(\frac{\int_{\mathcal{X}} p(x)q(x)dx}{\left(\int_{\mathcal{X}} p(x)^\alpha dx\right)^{\frac{1}{\alpha}} \left(\int_{\mathcal{X}} q(x)^\beta dx\right)^{\frac{1}{\beta}}} \right). \quad (10)$$

When $0 < \alpha < 1$ and $\beta = \bar{\alpha} = \frac{\alpha}{\alpha-1} < 0$, or $\alpha < 0$ and $0 < \beta < 1$, the reverse HPD is defined by:

$$D_\alpha^H(p(x) : q(x)) = \log \left(\frac{\int_{\mathcal{X}} p(x)q(x)dx}{\left(\int_{\mathcal{X}} p(x)^\alpha dx\right)^{\frac{1}{\alpha}} \left(\int_{\mathcal{X}} q(x)^\beta dx\right)^{\frac{1}{\beta}}} \right). \quad (11)$$

By Hölder's inequality and the reverse Hölder inequality, $D_\alpha^H(p(x) : q(x)) \geq 0$ with $D_\alpha^H(p(x) : q(x)) = 0$ iff $p(x)^\alpha \propto q(x)^\beta$ or equivalently $q(x) \propto p(x)^{\alpha/\beta} = p(x)^{\alpha-1}$. When $\alpha > 1$, $x^{\alpha-1}$ is monotonically increasing, and D_α^H is indeed a pseudo-divergence. However, the reverse HPD is *not* a pseudo-divergence because $x^{\alpha-1}$ will be monotonically decreasing if $\alpha < 0$ or $0 < \alpha < 1$. Therefore we only consider HPD with $\alpha > 1$ in the remainder, and leave here the notion of reverse Hölder divergence.

When $\alpha = \beta = 2$, the HPD becomes the Cauchy-Schwarz divergence CS [16]:

$$D_2^H(p(x) : q(x)) = \text{CS}(p(x) : q(x)) = -\log \left(\frac{\int_{\mathcal{X}} p(x)q(x)dx}{\left(\int_{\mathcal{X}} p(x)^2 dx\right)^{\frac{1}{2}} \left(\int_{\mathcal{X}} q(x)^2 dx\right)^{\frac{1}{2}}} \right), \quad (12)$$

which has been proved useful to get closed-form divergence formulas between mixtures of exponential families with conic or affine natural parameter spaces [21].

The Cauchy-Schwarz divergence is proper for probability densities since the Cauchy-Schwarz inequality becomes an equality iff $q(x) = \lambda p(x)^{\alpha-1} = \lambda p(x)$ implies that $\lambda = \int_{\mathcal{X}} \lambda p(x)dx = \int_{\mathcal{X}} q(x)dx = 1$. It is however not proper for positive densities.

Fact 1 (CS is only proper for probability densities). *The Cauchy-Schwarz divergence $\text{CS}(p(x) : q(x))$ is proper for square-integrable probability densities $p(x), q(x) \in L^2(\mathcal{X}, \mu)$ but not proper for positive square-integrable densities.*

2.2 Properness and improperness

In the general case, when $\alpha \neq 2$, the divergence D_α^H is not even proper for normalized (probability) densities, not to mention general unnormalized (positive) densities. Indeed, when $p(x) = q(x)$, we have:

$$D_\alpha^H(p(x) : p(x)) = -\log \left(\frac{\int p(x)^2 dx}{\left(\int p(x)^\alpha dx\right)^{\frac{1}{\alpha}} \left(\int p(x)^{\frac{\alpha}{\alpha-1}} dx\right)^{\frac{\alpha-1}{\alpha}}} \right) \neq 0 \text{ when } \alpha \neq 2. \quad (13)$$

Let us consider the general case. For unnormalized positive distributions $\tilde{p}(x)$ and $\tilde{q}(x)$ (the tilde notation stems from the notation of homogeneous coordinates in projective geometry), the inequality becomes an equality when: $\tilde{p}(x)^\alpha \propto \tilde{q}(x)^\beta$, i.e., $p(x)^\alpha \propto q(x)^\beta$, or $q(x) \propto p(x)^{\alpha/\bar{\alpha}} = p(x)^{\alpha-1}$. We can check that $D_\alpha^H(p(x) : \lambda p(x)^{\alpha-1}) = 0$ for any $\lambda > 0$:

$$-\log \left(\frac{\int p(x)\lambda p(x)^{\alpha-1} dx}{\left(\int p(x)^\alpha dx\right)^{\frac{1}{\alpha}} \left(\int \lambda^\beta p(x)^{(\alpha-1)\beta} dx\right)^{\frac{1}{\beta}}} \right) = -\log \left(\frac{\int p(x)^\alpha dx}{\left(\int p(x)^\alpha dx\right)^{\frac{1}{\alpha}} \left(\int p(x)^\alpha dx\right)^{\frac{1}{\beta}}} \right) = 0, \quad (14)$$

since $(\alpha-1)\beta = (\alpha-1)\bar{\alpha} = (\alpha-1)\frac{\alpha}{\alpha-1} = \alpha$.

For $\alpha = 2$, we find indeed that $D_2^H(p(x) : \lambda p(x)) = \text{CS}(p(x) : p(x)) = 0$ for any $\lambda \neq 0$.

Fact 2 (HPD is improper). *The Hölder pseudo-divergences are improper statistical distances.*

2.3 Reference duality

In general, Hölder divergences are asymmetric when $\alpha \neq \beta$ ($\neq 2$) but enjoy the following *reference duality* [39]:

$$D_\alpha^H(p(x) : q(x)) = D_\beta^H(q(x) : p(x)) = D_{\frac{\alpha}{\alpha-1}}^H(q(x) : p(x)). \quad (15)$$

Fact 3 (Reference duality HPD). *The Hölder pseudo-divergences satisfy the reference duality $\alpha \leftrightarrow \beta = \frac{\alpha}{\alpha-1}$: $D_\alpha^H(p(x) : q(x)) = D_\beta^H(q(x) : p(x)) = D_{\frac{\alpha}{\alpha-1}}^H(q(x) : p(x))$.*

An arithmetic symmetrization of the HPD yields a symmetric HPD S_α^H , given by:

$$\begin{aligned} S_\alpha^H(p(x) : q(x)) &= S_\alpha^H(q(x) : p(x)) = \frac{D_\alpha^H(p(x) : q(x)) + D_\alpha^H(q(x) : p(x))}{2}, \\ &= -\log \left(\frac{\int p(x)q(x)dx}{\sqrt{(\int p(x)^\alpha dx)^{\frac{1}{\alpha}} (\int p(x)^{\bar{\alpha}} dx)^{\frac{1}{\bar{\alpha}}} (\int q(x)^\alpha dx)^{\frac{1}{\alpha}} (\int q(x)^{\bar{\alpha}} dx)^{\frac{1}{\bar{\alpha}}}}} \right). \end{aligned} \quad (16)$$

2.4 HPD is a projective divergence

In the above definition, densities $p(x)$ and $q(x)$ can either be positive or normalized probability distributions. Let $\tilde{p}(x)$ and $\tilde{q}(x)$ denote positive (not necessarily normalized) measures, and $w(\tilde{p}) = \int_{\mathcal{X}} \tilde{p}(x)dx$ the *overall mass* so that $p(x) = \frac{\tilde{p}(x)}{w(\tilde{p})}$ is the corresponding normalized probability measure. Then we check that HPD is a *projective divergence* [13] since:

$$D_\alpha^H(\tilde{p}(x) : \tilde{q}(x)) = D_\alpha^H(p(x) : q(x)), \quad (17)$$

or in general:

$$D_\alpha^H(\lambda p(x) : \lambda' q(x)) = D_\alpha^H(p(x) : q(x)) \quad (18)$$

for all prescribed constants $\lambda, \lambda' > 0$. Projective divergences may also be called “*angular divergences*” or “*cosine divergences*” since they do not depend on the total mass of the measure densities.

Fact 4 (HPD is projective). *The Hölder pseudo-divergences are projective distances.*

2.5 Escort distributions and skew Bhattacharyya divergences

Let us define with respect to the probability measures $p(x) \in L^{\frac{1}{\alpha}}(\mathcal{X}, \mu)$ and $q(x) \in L^{\frac{1}{\beta}}(\mathcal{X}, \mu)$ the following *escort probability distributions* [2]:

$$p_\alpha^E(x) = \frac{p(x)^{\frac{1}{\alpha}}}{\int p(x)^{\frac{1}{\alpha}} dx}, \quad (19)$$

and

$$q_\beta^E(x) = \frac{q(x)^{\frac{1}{\beta}}}{\int q(x)^{\frac{1}{\beta}} dx}. \quad (20)$$

Since HPD is a projective divergence, we compute with respect to the conjugate exponents α and β the *Hölder Escort Divergence* (HED):

$$D_\alpha^{\text{HE}}(p(x) : q(x)) = D_\alpha^H(p_\alpha^E(x) : q_\beta^E(x)) = -\log \int_{\mathcal{X}} p(x)^{1/\alpha} q(x)^{1-1/\alpha} dx = B_{1/\alpha}(p(x) : q(x)), \quad (21)$$

which turns out to be the familiar *skew Bhattacharyya divergence* $B_{1/\alpha}(p(x) : q(x))$, see [22].

Fact 5 (HED as a skew Bhattacharyya divergence). *The Hölder escort divergence amounts to a skew Bhattacharyya divergence: $D_\alpha^{\text{HE}}(p(x) : q(x)) = B_{1/\alpha}(p(x) : q(x))$ for any $\alpha > 0$.*

In particular, the *Cauchy-Schwarz escort divergence* $\text{CS}^{\text{HE}}(p(x) : q(x))$ amounts to the Bhattacharyya distance [6] $B(p(x) : q(x)) = -\log \sqrt{\int_{\mathcal{X}} p(x)q(x)dx}$:

$$\text{CS}^{\text{HE}}(p(x) : q(x)) = D_2^{\text{HE}}(p(x) : q(x)) = D_2^H(p_2^E(x) : q_2^E(x)) = B_{1/2}(p(x) : q(x)) = B(p(x) : q(x)). \quad (22)$$

Observe that the Cauchy-Schwarz escort distributions are the square root density representations [36] of distributions.

3 Proper Hölder divergence

3.1 Definition

Let $p(x)$ and $q(x)$ be positive measures in $L^\gamma(\mathcal{X}, \mu)$ for a prescribed scalar value $\gamma > 0$. Plugging the positive measures $p(x)^{\gamma/\alpha}$ and $q(x)^{\gamma/\beta}$ into the definition of HPD D_α^H , we get the following definition:

Definition 2 (Proper Hölder Divergence, HD). *For conjugate exponents $\alpha, \beta > 0$ and $\gamma > 0$, the proper Hölder divergence between two densities $p(x)$ and $q(x)$ is defined by:*

$$D_{\alpha,\gamma}^H(p(x) : q(x)) = D_\alpha^H(p(x)^{\gamma/\alpha} : q(x)^{\gamma/\beta}) = -\log \left(\frac{\int_{\mathcal{X}} p(x)^{\gamma/\alpha} q(x)^{\gamma/\beta} dx}{(\int_{\mathcal{X}} p(x)^{\gamma} dx)^{1/\alpha} (\int_{\mathcal{X}} q(x)^{\gamma} dx)^{1/\beta}} \right). \quad (23)$$

By definition, $D_{\alpha,\gamma}^H(p : q)$ is a *two-parameter family* of dissimilarity statistical measures. Following Hölder's inequality, we can check that $D_{\alpha,\gamma}^H(p(x) : q(x)) \geq 0$ and $D_{\alpha,\gamma}^H(p(x) : q(x)) = 0$ iff $p(x)^\gamma \propto q(x)^\gamma$, i.e. $p(x) \propto q(x)$ (see Appendix A). If $p(x)$ and $q(x)$ belong to the statistical probability manifold, then $D_{\alpha,\gamma}^H(p(x) : q(x)) = 0$ iff $p(x) = q(x)$ almost everywhere. This says that HD is a proper divergence for probability measures, and it becomes a pseudo-divergence for positive measures. Note that we have abused the notation D^H to denote both the Hölder pseudo-divergence (with one subscript) and the Hölder divergence (with two subscripts).

Similar to HPD, HD is asymmetric when $\alpha \neq \beta$ with the following reference duality:

$$D_{\alpha,\gamma}^H(p(x) : q(x)) = D_{\alpha,\gamma}^H(q(x), p(x)). \quad (24)$$

HD can be symmetrized as:

$$S_{\alpha,\gamma}^H(p : q) = \frac{D_{\alpha,\gamma}^H(p : q) + D_{\alpha,\gamma}^H(q : p)}{2} = -\log \sqrt{\frac{\int_{\mathcal{X}} p(x)^{\gamma/\alpha} q(x)^{\gamma/\beta} dx \int_{\mathcal{X}} p(x)^{\gamma/\beta} q(x)^{\gamma/\alpha} dx}{\int_{\mathcal{X}} p(x)^{\gamma} dx \int_{\mathcal{X}} q(x)^{\gamma} dx}}. \quad (25)$$

Furthermore, one can easily check that HD is a projective divergence.

For conjugate exponents $\alpha, \beta > 0$ and $\gamma > 0$, we rewrite the definition of HD as:

$$\begin{aligned} D_{\alpha,\gamma}^H(p(x) : q(x)) &= -\log \int_{\mathcal{X}} \left(\frac{p(x)^\gamma}{\int_{\mathcal{X}} p(x)^\gamma dx} \right)^{1/\alpha} \left(\frac{q(x)^\gamma}{\int_{\mathcal{X}} q(x)^\gamma dx} \right)^{1/\beta} dx, \\ &= -\log \left(p_{1/\gamma}^E(x) \right)^{1/\alpha} \left(q_{1/\gamma}^E(x) \right)^{1/\beta} dx = B_{\frac{1}{\alpha}}(p_{1/\gamma}^E(x) : q_{1/\gamma}^E(x)). \end{aligned}$$

Therefore HD can be reinterpreted as the skew Bhattacharyya divergence [22] between the escort distributions. In particular, when $\gamma = 1$, we get:

$$D_{\alpha,1}^H(p(x) : q(x)) = -\log \left(\int_{\mathcal{X}} p(x)^{1/\alpha} q(x)^{1/\beta} dx \right) = B_{\frac{1}{\alpha}}(p(x) : q(x)). \quad (26)$$

Fact 6. *The two-parametric family of statistical Hölder divergence $D_{\alpha,\gamma}^H$ passes through the one-parametric family of skew Bhattacharyya divergences when $\gamma = 1$.*

3.2 Special case: The Cauchy-Schwarz divergence

We consider the intersection of the uni-parametric class of Hölder pseudo-divergences (HPD) with the bi-parametric class of proper Hölder divergences (HD): That is, the class of divergences which belong to both HPD and HD. Then we must have $\gamma/\alpha = \gamma/\beta = 1$. Since $1/\alpha + 1/\beta = 1$, we get $\alpha = \beta = \gamma = 2$. Therefore the Cauchy-Schwarz (CS) divergence is the *unique* divergence belonging to both HPD and HD classes:

$$D_{2,2}^H(p(x) : q(x)) = D_2^H(p(x) : q(x)) = \text{CS}(p(x) : q(x)). \quad (27)$$

In fact, the CS divergence is the intersection of the four classes HPD, symmetric HPD, HD, and symmetric HD. Figure 1 displays a diagram of those divergence classes with their inclusion relationships.

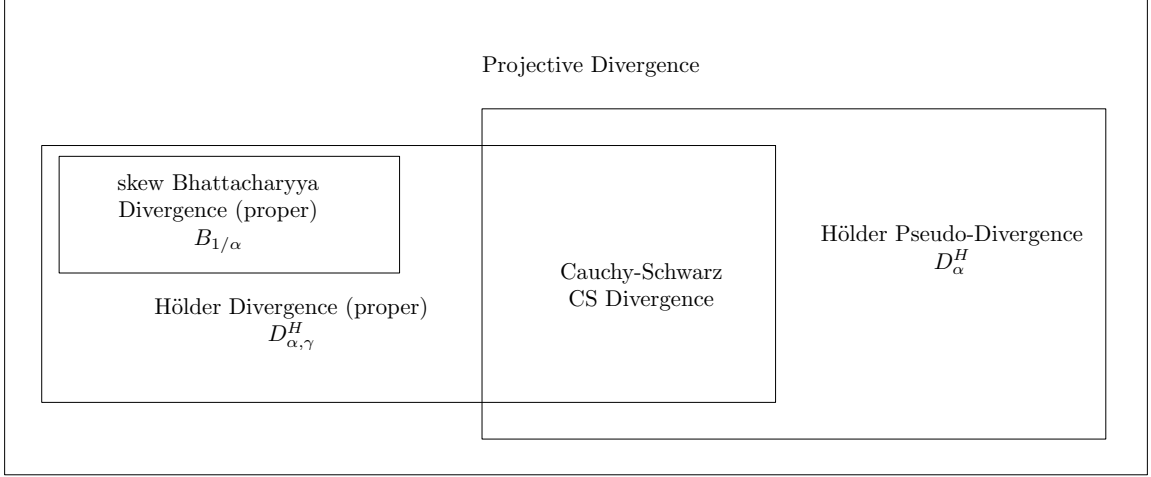


Figure 1: Hölder proper divergence (bi-parametric) and Hölder improper pseudo-divergence (uni-parametric) intersect at the unique non-parametric Cauchy-Schwarz divergence. By using escort distributions, Hölder divergences encapsulates the skew Bhattacharyya distances.

As stated earlier, notice that the Cauchy-Schwarz inequality:

$$\int p(x)q(x)dx \leq \sqrt{\left(\int p(x)^2 dx\right) \left(\int q(x)^2 dx\right)}, \quad (28)$$

is not proper as it is an equality when $p(x)$ and $q(x)$ are *linearly dependent* (i.e., $p(x) = \lambda q(x)$ for $\lambda > 0$). The arguments of the CS divergence are square-integrable real-valued density functions $p(x)$ and $q(x)$. Thus the Cauchy-Schwarz divergence is not proper for positive measures but is proper for normalized probability distributions since $\int p(x)dx = \int \lambda q(x)dx = 1$ implies that $\lambda = 1$.

3.3 Limit cases of Hölder divergences and statistical estimation

Let us define the inner product of unnormalized densities as:

$$\langle \tilde{p}(x), \tilde{q}(x) \rangle = \int_{\mathcal{X}} \tilde{p}(x)\tilde{q}(x)dx \quad (29)$$

(for $L^2(\mathcal{X}, \mu)$ integrable functions), and define the L_α norm of densities as $\|\tilde{p}(x)\|_\alpha = (\int_{\mathcal{X}} \tilde{p}(x)^\alpha dx)^{1/\alpha}$ for $\alpha \geq 1$. Then the CS divergence can be concisely written as:

$$\text{CS}(\tilde{p}(x), \tilde{q}(x)) = -\log \frac{\langle \tilde{p}(x), \tilde{q}(x) \rangle}{\|\tilde{p}(x)\|_2 \|\tilde{q}(x)\|_2}, \quad (30)$$

and the Hölder pseudo-divergence writes as:

$$D_\alpha^H(\tilde{p}(x), \tilde{q}(x)) = -\log \frac{\langle \tilde{p}(x), \tilde{q}(x) \rangle}{\|\tilde{p}(x)\|_\alpha \|\tilde{q}(x)\|_{\bar{\alpha}}}. \quad (31)$$

When $\alpha \rightarrow 1^+$, we have $\bar{\alpha} = \alpha/(\alpha - 1) \rightarrow +\infty$. Then it comes that:

$$\lim_{\alpha \rightarrow 1^+} D_\alpha^H(\tilde{p}(x), \tilde{q}(x)) = -\log \frac{\langle \tilde{p}(x), \tilde{q}(x) \rangle}{\|\tilde{p}(x)\|_1 \|\tilde{q}(x)\|_\infty} = -\log \langle \tilde{p}(x), \tilde{q}(x) \rangle + \log \int_{\mathcal{X}} \tilde{p}(x)dx + \log \max_{x \in \mathcal{X}} \tilde{q}(x). \quad (32)$$

When $\alpha \rightarrow +\infty$ and $\bar{\alpha} \rightarrow 1^+$, we have:

$$\lim_{\alpha \rightarrow +\infty} D_\alpha^H(\tilde{p}(x), \tilde{q}(x)) = -\log \frac{\langle \tilde{p}(x), \tilde{q}(x) \rangle}{\|\tilde{p}(x)\|_\infty \|\tilde{q}(x)\|_1} = -\log \langle \tilde{p}(x), \tilde{q}(x) \rangle + \log \max_{x \in \mathcal{X}} \tilde{p}(x) + \log \int_{\mathcal{X}} \tilde{q}(x)dx. \quad (33)$$

Now consider a pair of probability densities $p(x)$ and $q(x)$. We have:

$$\begin{aligned}\lim_{\alpha \rightarrow 1^+} D_\alpha^H(p(x), q(x)) &= -\log \langle p(x), q(x) \rangle + \max_{x \in \mathcal{X}} \log q(x), \\ \lim_{\alpha \rightarrow +\infty} D_\alpha^H(p, q) &= -\log \langle p(x), q(x) \rangle + \max_{x \in \mathcal{X}} \log p(x), \\ D_2^H(p, q) &= -\log \langle p(x), q(x) \rangle + \log \|p(x)\|_2 + \log \|q(x)\|_2.\end{aligned}\quad (34)$$

In an estimation scenario, $p(x)$ is fixed and $q(x|\theta) = q_\theta(x)$ is free along a parametric manifold \mathcal{M} , then minimizing Hölder divergence reduces to:

$$\begin{aligned}\arg \min_{\theta \in \mathcal{M}} \lim_{\alpha \rightarrow 1^+} D_\alpha^H(p(x), q_\theta(x)) &= \arg \min_{\theta \in \mathcal{M}} \left(-\log \langle p(x), q_\theta(x) \rangle + \max_{x \in \mathcal{X}} \log q_\theta(x) \right), \\ \arg \min_{\theta \in \mathcal{M}} \lim_{\alpha \rightarrow +\infty} D_\alpha^H(p(x), q(x)) &= \arg \min_{\theta \in \mathcal{M}} \left(-\log \langle p(x), q_\theta(x) \rangle \right), \\ \arg \min_{\theta \in \mathcal{M}} D_2^H(p(x), q(x)) &= \arg \min_{\theta \in \mathcal{M}} \left(-\log \langle p(x), q_\theta(x) \rangle + \log \|q_\theta(x)\|_2 \right).\end{aligned}\quad (35)$$

Therefore when θ varies from 1 to $+\infty$, only the regularizer in the minimization problem changes. In any case, Hölder divergence always has the term $-\log \langle p(x), q(x) \rangle$, which shares a similar form with the Bhattacharyya distance [6]:

$$B(p(x) : q(x)) = -\log \int_{\mathcal{X}} \sqrt{p(x)q(x)} dx = -\log \langle \sqrt{p(x)}, \sqrt{q(x)} \rangle. \quad (36)$$

HPD between $\tilde{p}(x)$ and $\tilde{q}(x)$ is also closely related to their cosine similarity $\frac{\langle \tilde{p}(x), \tilde{q}(x) \rangle}{\|\tilde{p}(x)\|_2 \|\tilde{q}(x)\|_2}$. When $\alpha = 2$, HD is exactly the cosine similarity after a non-linear transformation.

4 Closed-form expressions of HPD and HD for conic and affine exponential families

We report closed-form formulas for the HPD and HD between two distributions belonging to the same exponential family provided that the natural parameter space is a cone or affine. A cone Ω is a convex domain such that for $P, Q \in \Omega$ and any $\lambda > 0$, we have $P + \lambda Q \in \Omega$. For example, the set of positive measures absolutely continuous with a base measure μ is a cone. Recall that an exponential family [23] has a density function $p(x; \theta)$ that we be written canonically as:

$$p(x; \theta) = \exp(\langle t(x), \theta \rangle - F(\theta) + k(x)). \quad (37)$$

In this work, we consider the auxiliary carrier measure term $k(x) = 0$. The base measure is either the Lebesgue measure μ or the counting measure μ_C . A Conic or Affine Exponential Family (CAEF) is an exponential family with the natural parameter space Θ a cone or affine. The log-normalizer $F(\theta)$ is a strictly convex function also called cumulant generating function [2].

Lemma 1 (HPD and HD for CAEFs). *For distributions $p(x; \theta_p)$ and $p(x; \theta_q)$ belonging to the same exponential family with conic or affine natural parameter space [21], both the HPD and HD are available in closed-form:*

$$D_\alpha^H(p : q) = \frac{1}{\alpha} F(\alpha \theta_p) + \frac{1}{\beta} F(\beta \theta_q) - F(\theta_p + \theta_q), \quad (38)$$

$$D_{\alpha, \gamma}^H(p : q) = \frac{1}{\alpha} F(\gamma \theta_p) + \frac{1}{\beta} F(\gamma \theta_q) - F\left(\frac{\gamma}{\alpha} \theta_p + \frac{\gamma}{\beta} \theta_q\right). \quad (39)$$

Proof. Consider $k(x) = 0$ and a conic or affine natural space Θ (see [21]), then for all $a, b > 0$, we have:

$$\left(\int p(x)^a dx \right)^{\frac{1}{b}} = \exp\left(\frac{1}{b} F(a \theta_p) - \frac{a}{b} F(\theta_p)\right), \quad (40)$$

since $a\theta_p \in \Theta$. Indeed, we have:

$$\begin{aligned} \left(\int p(x)^a dx \right)^{1/b} &= \left(\int \exp(\langle a\theta, t(x) \rangle - aF(\theta)) dx \right)^{1/b} \\ &= \left(\int \exp(\langle a\theta, t(x) \rangle - F(a\theta) + F(a\theta) - aF(\theta)) dx \right)^{1/b} \\ &= \exp\left(\frac{1}{b}F(a\theta) - \frac{a}{b}F(\theta)\right) \underbrace{\left(\int \exp(\langle a\theta, t(x) \rangle - F(a\theta)) dx \right)}_{=1}^{1/b}. \end{aligned}$$

Similarly, we have for all $a, b > 0$ (details omitted),

$$\int p(x)^a q(x)^b dx = \exp(F(a\theta_p + b\theta_q) - aF(\theta_p) - bF(\theta_q)), \quad (41)$$

since $a\theta_p + b\theta_q \in \Theta$. Therefore, we get:

$$D_\alpha^H(p(x) : q(x)) = -\log \frac{\int p(x)q(x)dx}{\left(\int p(x)^\alpha dx\right)^{\frac{1}{\alpha}} \left(\int q(x)^\beta dx\right)^{\frac{1}{\beta}}} \quad (42)$$

$$= -F(\theta_p + \theta_q) + F(\theta_p) + F(\theta_q) + \frac{1}{\alpha}F(\alpha\theta_p) - F(\theta_p) + \frac{1}{\beta}F(\beta\theta_q) - F(\theta_q) \quad (43)$$

$$= \frac{1}{\alpha}F(\alpha\theta_p) + \frac{1}{\beta}F(\beta\theta_q) - F(\theta_p + \theta_q) \geq 0, \quad (44)$$

$$D_{\alpha,\gamma}^H(p(x) : q(x)) = -\log \frac{\int p(x)^{\gamma/\alpha} q(x)^{\gamma/\beta} dx}{\left(\int p(x)^\gamma dx\right)^{\frac{1}{\alpha}} \left(\int q(x)^\gamma dx\right)^{\frac{1}{\beta}}} \quad (45)$$

$$= -F\left(\frac{\gamma}{\alpha}\theta_p + \frac{\gamma}{\beta}\theta_q\right) + \frac{\gamma}{\alpha}F(\theta_p) + \frac{\gamma}{\beta}F(\theta_q) + \frac{1}{\alpha}F(\gamma\theta_p) - \frac{\gamma}{\alpha}F(\theta_p) + \frac{1}{\beta}F(\gamma\theta_q) - \frac{\gamma}{\beta}F(\theta_q) \quad (46)$$

$$= \frac{1}{\alpha}F(\gamma\theta_p) + \frac{1}{\beta}F(\gamma\theta_q) - F\left(\frac{\gamma}{\alpha}\theta_p + \frac{\gamma}{\beta}\theta_q\right) \geq 0. \quad (47)$$

When $1 > \alpha > 0$, we have $\beta = \frac{\alpha}{\alpha-1} < 0$. To get similar results for the reverse Hölder divergence, we need the natural parameter space to be affine (eg., isotropic Gaussians or multinomials, see [27]). \square

In particular, if $p(x)$ and $q(x)$ belong to the same exponential family so that $p(x) = \exp(\langle \theta_p, t(x) \rangle - F(\theta_p))$ and $q(x) = \exp(\langle \theta_q, t(x) \rangle - F(\theta_q))$, one can easily check that $D_\alpha^H(p(x; \theta_p) : p(x; \theta_q)) = 0$ iff $\theta_q = (\alpha - 1)\theta_p$. For HD, we can check $D_{\alpha,\gamma}^H(p(x) : p(x)) = 0$ is proper since $\frac{1}{\alpha} + \frac{1}{\beta} = 1$.

The following result is straightforward from Lemma 1.

Lemma 2 (Symmetric HPD and HD for CAEFs). *For distributions $p(x; \theta_p)$ and $p(x; \theta_q)$ belonging to the same exponential family with conic or affine natural parameter space [21], the symmetric HPD and HD are available in closed-form:*

$$S_\alpha^H(p(x) : q(x)) = \frac{1}{2} \left[\frac{1}{\alpha}F(\alpha\theta_p) + \frac{1}{\beta}F(\beta\theta_p) + \frac{1}{\alpha}F(\alpha\theta_q) + \frac{1}{\beta}F(\beta\theta_q) \right] - F(\theta_p + \theta_q); \quad (48)$$

$$S_{\alpha,\gamma}^H(p(x) : q(x)) = \frac{1}{2} \left[F(\gamma\theta_p) + F(\gamma\theta_q) - F\left(\frac{\gamma}{\alpha}\theta_p + \frac{\gamma}{\beta}\theta_q\right) - F\left(\frac{\gamma}{\beta}\theta_p + \frac{\gamma}{\alpha}\theta_q\right) \right]. \quad (49)$$

Remark 1. *By reference duality,*

$$\begin{aligned} S_\alpha^H(p(x) : q(x)) &= S_\alpha^H(p(x) : q(x)); \\ S_{\alpha,\gamma}^H(p(x) : q(x)) &= S_{\alpha,\gamma}^H(p(x) : q(x)). \end{aligned}$$

Note that the Hölder score-induced divergence [19] does *not* admit in general closed-form formulas for exponential families since it relies on a *function* $\phi(\cdot)$ (see Definition 4 of [19]).

Note that CAEF convex log-normalizers satisfy:

$$\frac{1}{\alpha}F(\alpha\theta_p) + \frac{1}{\beta}F(\beta\theta_q) \geq F(\theta_p + \theta_q). \quad (50)$$

A necessary condition is that $F(\lambda\theta) \geq \lambda F(\theta)$ for $\lambda > 0$ (take $\theta_p = \theta$, $\theta_q = 0$ and $F(0) = 0$ in the above equality).

The escort distribution for an exponential family is given by:

$$p_\alpha^E(x; \theta) = e^{\frac{F(\theta)}{\alpha} - F(\frac{\theta}{\alpha})} p(x; \theta)^{\frac{1}{\alpha}}. \quad (51)$$

The Hölder equality holds when $p(x)^\alpha \propto q(x)^\beta$ or $p(x)^\alpha q(x)^{-\beta} \propto 1$. For exponential families, this condition is satisfied when $\alpha\theta_p - \beta\theta_q \in \Theta$. That is, we need to have:

$$\alpha \left(\theta_p - \frac{1}{\alpha - 1} \theta_q \right) \in \Theta. \quad (52)$$

Thus we may choose small enough $\alpha = 1 + \epsilon > 1$ so that the condition is not satisfied for fixed θ_p and θ_q for many exponential distributions. Since multinomials have affine natural space [27], this condition is always met, but not for non-affine natural parameter spaces like normal distributions.

Notice the following fact:

Fact 7 (Density of a CAEF in $L^\gamma(\mathcal{X}, \mu)$). *The density of exponential families with conic or affine natural parameter space belongs to $L^\gamma(\mathcal{X}, \mu)$ for any $\gamma > 0$.*

Proof. We have $\int_{\mathcal{X}} (\exp(\langle \theta, t(x) \rangle) - F(\theta))^\gamma d\mu(x) = e^{F(\gamma\theta) - \gamma F(\theta)} < \infty$ for any $\gamma > 0$ provided that $\gamma\theta$ belongs to the natural parameter space. When Θ is a cone or affine, the condition is satisfied. \square

Let $\tilde{p}(x; \theta) = \exp(\langle t(x), \theta \rangle)$ denote the unnormalized positive exponential family density and $p(x; \theta) = \frac{\tilde{p}(x; \theta)}{Z(\theta)}$ the normalized density with $Z(\theta) = \exp(F(\theta))$ the partition function. Although HD is a projective divergence since we have $D_\alpha^H(p(x; \theta_1) : p(x; \theta_2)) = D_\alpha^H(\tilde{p}(x; \theta_1) : \tilde{p}(x; \theta_2))$, observe that the HD value *depends* on the log-normalizer $F(\theta)$ (since the HD is an integral on the support, see also [29] for a similar argument with the γ -divergence [13]).

In practice, even when the log-normalizer is computationally intractable, we may still estimate the HD by Monte-Carlo techniques: Indeed, we can sample a distribution $\tilde{p}(x)$ either by rejection sampling [29] or by the Markov Chain Monte-Carlo (MCMC) Metropolis-Hasting technique: It just requires to be able to sample a proposal distribution that has the same support.

We shall now instantiate the HPD and HD formulas for several exponential families with conic or affine natural parameter spaces.

4.1 Case study: Categorical distributions

Let $p = (p_0, \dots, p_m)$ and $q = (q_0, \dots, q_m)$ be two categorical distributions in the m -dimensional probability simplex Δ^m . We rewrite p in the canonical form of exponential families [23] as:


$$p_i = \exp \left((\theta_p)_i - \log \left(1 + \sum_{i=1}^m \exp(\theta_p)_i \right) \right), \quad \forall i \in \{1, \dots, m\}, \quad (53)$$

with the redundant parameter:

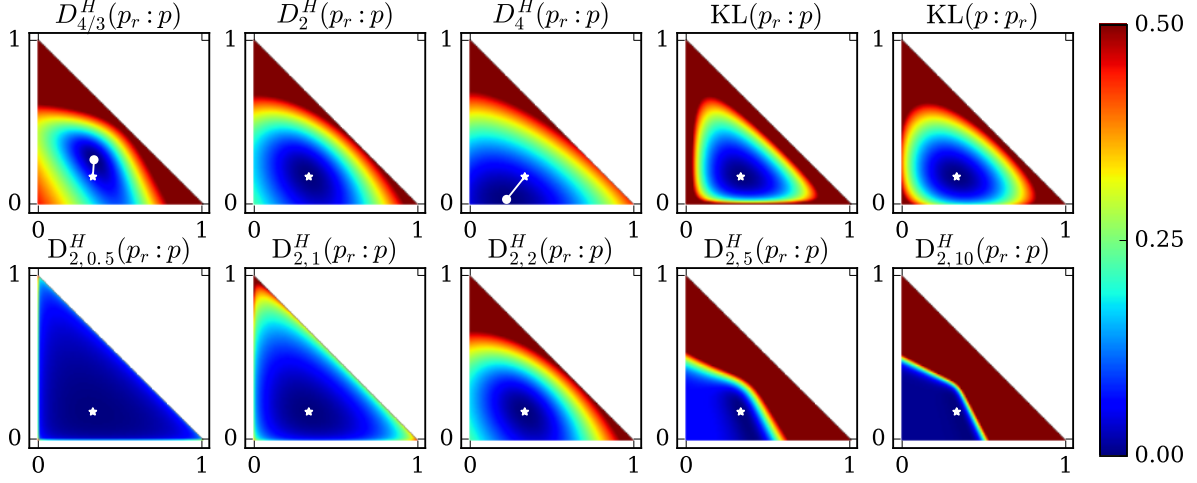
$$p_0 = 1 - \sum_{i=1}^m p_i = \frac{1}{1 + \sum_{i=1}^m \exp(\theta_p)_i}. \quad (54)$$

From Eq. 53, the convex cumulant generating function has the form $F(\theta) = \log(1 + \sum_{i=1}^m \exp(\theta_p)_i)$. The inverse transformation from p to θ is therefore given by:

$$\theta_i = \log \left(\frac{p_i}{p_0} \right), \quad \forall i \in \{1, \dots, m\}. \quad (55)$$



(a) Reference categorical distribution $p_r = (1/3, 1/3, 1/3)$



(b) Reference categorical distribution $p_r = (1/2, 1/3, 1/6)$

Figure 2: $D_{\alpha}^H(p_r : p)$ for $\alpha \in \{4, 2, 4/3\}$, and $D_{2,\gamma}^H(p_r : p)$ for $\gamma \in \{0.5, 1, 2, 5, 10\}$, compared to KL divergence. The reference distribution p_r is presented as "★". The minimizer of $D_{\alpha}^H(p_0 : p)$, if different from p_0 , is presented as "●". Notice that $D_{2,2}^H = D_2^H$.

The natural parameter space Θ is affine (hence conic), and by applying Lemma 1, we get the following closed-form formula:

$$D_{\alpha}^H(p : q) = \frac{1}{\alpha} \log \left(1 + \sum_{i=1}^m \exp(\alpha(\theta_p)_i) \right) + \frac{1}{\beta} \log \left(1 + \sum_{i=1}^m \exp(\beta(\theta_q)_i) \right) - \log \left(1 + \sum_{i=1}^m \exp((\theta_p)_i + (\theta_q)_i) \right), \quad (56)$$

$$D_{\alpha,\gamma}^H(p : q) = \frac{1}{\alpha} \log \left(1 + \sum_{i=1}^m \exp(\gamma(\theta_p)_i) \right) + \frac{1}{\beta} \log \left(1 + \sum_{i=1}^m \exp(\gamma(\theta_q)_i) \right) - \log \left(1 + \sum_{i=1}^m \exp \left(\frac{\gamma}{\alpha}(\theta_p)_i + \frac{\gamma}{\beta}(\theta_q)_i \right) \right). \quad (57)$$

To get some intuitions, Fig. 2 shows the Hölder divergence from any categorical distribution (p_0, p_1, p_2) to a given reference distribution p_r in the 2D probability simplex Δ^2 . A main observation is that the

Kullback-Leibler (KL) divergence exhibits a barrier near the boundary $\partial\Delta^2$ with large values. This is *not* the case for Hölder divergences: $D_\alpha^H(p : p_r)$ does not have a sharp increase near the boundary (although it still penalizes the corners of Δ^2). For example, let $p = (0, 1/2, 1/2)$, $p_r = (1/3, 1/3, 1/3)$, then $\text{KL}(p_r : p) \rightarrow \infty$ but $D_2^H(p_r : p) = \sqrt{2/3}$. Another observation is that the minimum $D(p : p_r)$ can be reached at some point $p \neq p_r$ (see for example $D_4^H(p : p_r)$ in Fig. 2b, the bluest area corresponding to the minimum of $D(p : p_r)$ is not in the same location as the reference point).

Consider a HPD ball of center c and prescribed radius r wrt the HPD. Since $p(x)^{\alpha-1}$ for $\alpha \neq 2$ does not belong to the probability manifold but to the positive measure manifold and since the distance is projective, we deduce that the displaced ball center c' of a ball c lying the probability manifold can be computed as the intersection of the ray $0p^{\alpha-1}$ with the probability manifold. For the discrete probability simplex Δ , since we have $\lambda \sum_{x \in \mathcal{X}} p(x)^{\alpha-1} = 1$, we deduce that the displaced ball center is at:

$$c' = \frac{c}{\sum_{x \in \mathcal{X}} p(x)^{\alpha-1}} \quad (58)$$

This center is displayed as “•” in Figure 2.

In general, the HPD bisector [26] between two distributions belonging to the same CAEF is defined by:

$$\frac{1}{\alpha}(F(\alpha\theta_1) - F(\alpha\theta_2)) = F(\theta_2 + \theta) - F(\theta_1 + \theta). \quad (59)$$

4.2 Case study: Bernoulli distribution

Bernoulli distribution is just a special case of the category distribution when the number of categories is 2 (i.e., $m = 1$). To be consistent with the previous section, we rewrite a Bernoulli distribution $p = (p_0, p_1)$ in the canonical form:

$$p_1 = \exp(\theta_p - \log(1 + \exp(\theta_p))) = \frac{\exp(\theta_p)}{1 + \exp(\theta_p)}, \quad (60)$$

so that

$$p_0 = \frac{1}{1 + \exp(\theta_p)}. \quad (61)$$

Then the cumulant generating function becomes $F(\theta_p) = \log(1 + \exp(\theta_p))$. By Lemma 1,

$$D_\alpha^H(p(x) : q(x)) = \frac{1}{\alpha} \log(1 + \exp(\alpha\theta_p)) + \frac{1}{\beta} \log(1 + \exp(\beta\theta_q)) - \log(1 + \exp(\theta_p + \theta_q)), \quad (62)$$

$$D_{\alpha,\gamma}^H(p(x) : q(x)) = \frac{1}{\alpha} \log(1 + \exp(\gamma\theta_p)) + \frac{1}{\beta} \log(1 + \exp(\gamma\theta_q)) - \log\left(1 + \exp\left(\frac{\gamma}{\alpha}\theta_p + \frac{\gamma}{\beta}\theta_q\right)\right). \quad (63)$$

4.3 Case study: MultiVariate Normal distributions (MVNs)

Let us now instantiate the formulas for multivariate normals (Gaussian distributions). We have the log-normalizer $F(\theta)$ expressed using the usual parameters as [25]:

$$F(\theta) = F(\mu(\theta), \Sigma(\theta)) = \frac{1}{2} \log(2\pi)^d |\Sigma| + \frac{1}{2} \mu^\top \Sigma^{-1} \mu. \quad (64)$$

Since

$$\theta = (\Sigma^{-1} \mu, -\frac{1}{2} \Sigma^{-1}) = (v, M), \quad \mu = -\frac{1}{2} M^{-1} v, \quad \Sigma = -\frac{1}{2} M^{-1}. \quad (65)$$

It follows that:

$$\theta_p + \theta_q = \theta_{p+q} = (v_p + v_q, M_p + M_q) = \left(\Sigma_p^{-1} \mu_p + \Sigma_q^{-1} \mu_q, -\frac{1}{2} \Sigma_p^{-1} - \frac{1}{2} \Sigma_q^{-1} \right). \quad (66)$$

Therefore, we have:

$$\mu_{p+q} = (\Sigma_p^{-1} + \Sigma_q^{-1})^{-1} (\Sigma_p^{-1} \mu_p + \Sigma_q^{-1} \mu_q), \quad \Sigma_{p+q} = (\Sigma_p^{-1} + \Sigma_q^{-1})^{-1} \quad (67)$$

We thus get the following closed-form formula for $p \sim N(\mu_p, \Sigma_p)$ and $q \sim N(\mu_q, \Sigma_q)$:




Figure 3: First row: $D_{\alpha}^H(p_r : p)$, where p_r is standard Gaussian distribution, and $\alpha \in \{4/3, 2, 4\}$ as compared to the KL divergence. Second row: $D_{2,\gamma}^H(p_r : p)$ for $\gamma \in \{0.5, 1, 2, 5, 10\}$. Notice that $D_{2,2}^H = D_2^H$.

$$D_{\alpha}^H(N(\mu_p, \Sigma_p) : N(\mu_q, \Sigma_q)) = \frac{1}{2\alpha} \log \left| \frac{\Sigma_p}{\alpha} \right| + \frac{1}{2} \mu_p^{\top} \Sigma_p^{-1} \mu_p + \frac{1}{2\beta} \log \left| \frac{\Sigma_q}{\beta} \right| + \frac{1}{2} \mu_q^{\top} \Sigma_q^{-1} \mu_q \\ + \frac{1}{2} \log |\Sigma_p^{-1} + \Sigma_q^{-1}| - \frac{1}{2} (\Sigma_p^{-1} \mu_p + \Sigma_q^{-1} \mu_q)^{\top} (\Sigma_p^{-1} + \Sigma_q^{-1})^{-1} (\Sigma_p^{-1} \mu_p + \Sigma_q^{-1} \mu_q); \quad (68)$$

$$D_{\alpha,\gamma}^H(N(\mu_p, \Sigma_p) : N(\mu_q, \Sigma_q)) = \frac{1}{2\alpha} \log \left| \frac{\Sigma_p}{\gamma} \right| + \frac{\gamma}{2\alpha} \mu_p^{\top} \Sigma_p^{-1} \mu_p + \frac{1}{2\beta} \log \left| \frac{\Sigma_q}{\gamma} \right| + \frac{\gamma}{2\beta} \mu_q^{\top} \Sigma_q^{-1} \mu_q \\ + \frac{1}{2} \log \left| \frac{\gamma}{\alpha} \Sigma_p^{-1} + \frac{\gamma}{\beta} \Sigma_q^{-1} \right| - \frac{1}{2} \left(\frac{\gamma}{\alpha} \Sigma_p^{-1} \mu_p + \frac{\gamma}{\beta} \Sigma_q^{-1} \mu_q \right)^{\top} \left(\frac{\gamma}{\alpha} \Sigma_p^{-1} + \frac{\gamma}{\beta} \Sigma_q^{-1} \right)^{-1} \left(\frac{\gamma}{\alpha} \Sigma_p^{-1} \mu_p + \frac{\gamma}{\beta} \Sigma_q^{-1} \mu_q \right). \quad (69)$$

Figure 3 shows HPD and HD for univariate Gaussian distributions as compared to the KL divergence.

4.4 Case study: Zero-centered Laplace distribution

The zero-centered Laplace distribution is defined on the support $(-\infty, \infty)$ with the pdf:

$$p(x; \sigma) = \frac{1}{2\sigma} \exp \left(-\frac{|x|}{\sigma} \right) = \exp \left(-\frac{|x|}{\sigma} - \log(2\sigma) \right). \quad (70)$$

We have $\theta = -\frac{1}{\sigma}$, $F(\theta) = \log(-\frac{2}{\theta})$. Therefore, it comes that:

$$D_{\alpha}^H(p(x) : q(x)) = \frac{1}{\alpha} \log \left(-\frac{2}{\alpha \theta_p} \right) + \frac{1}{\beta} \log \left(-\frac{2}{\beta \theta_q} \right) - \log \left(-\frac{2}{\theta_p + \theta_q} \right) \\ = \frac{1}{\alpha} \log \left(\frac{\sigma_p}{\alpha} \right) + \frac{1}{\beta} \log \left(\frac{\sigma_q}{\beta} \right) - \log \left(\frac{\sigma_p \sigma_q}{\sigma_p + \sigma_q} \right), \quad (71)$$

$$D_{\alpha,\gamma}^H(p(x) : q(x)) = \frac{1}{\alpha} \log \left(-\frac{2}{\gamma \theta_p} \right) + \frac{1}{\beta} \log \left(-\frac{2}{\gamma \theta_q} \right) - \log \left(-\frac{2}{\frac{\gamma}{\alpha} \theta_p + \frac{\gamma}{\beta} \theta_q} \right) \\ = \frac{1}{\alpha} \log \sigma_p + \frac{1}{\beta} \log \sigma_q - \log \left(\frac{\sigma_p \sigma_q}{\frac{1}{\beta} \sigma_p + \frac{1}{\alpha} \sigma_q} \right). \quad (72)$$

In this special case, $D_{\alpha,\gamma}^H(p(x) : q(x))$ does not vary with γ .

4.5 Case study: Wishart distribution

The Wishart distribution is defined on the $d \times d$ positive definite cone with the density:

$$p(X; n, S) = \frac{|X|^{\frac{n-d-1}{2}} \exp\left(-\frac{1}{2} \text{tr}(S^{-1}X)\right)}{2^{\frac{nd}{2}} |S|^{\frac{n}{2}} \Gamma_d\left(\frac{n}{2}\right)}, \quad (73)$$

where $n > d - 1$ is the degree of freedom, and $S \succ 0$ is a scale matrix. We rewrite it into the canonical form:

$$p(X; n, S) = \exp\left(-\frac{1}{2} \text{tr}(S^{-1}X) + \frac{n-d-1}{2} \log |X| - \frac{nd}{2} \log 2 - \frac{n}{2} \log |S| - \log \Gamma_d\left(\frac{n}{2}\right)\right). \quad (74)$$

We can see that $\theta = (\theta^1, \theta^2)$, $\theta^1 = -\frac{1}{2}S^{-1}$, $\theta^2 = \frac{n-d-1}{2}$, and

$$\begin{aligned} F(\theta) &= \frac{nd}{2} \log 2 + \frac{n}{2} \log |S| + \log \Gamma_d\left(\frac{n}{2}\right) \\ &= \left(\theta^2 + \frac{d+1}{2}\right) d \log 2 + \left(\theta^2 + \frac{d+1}{2}\right) \log \left| -\frac{1}{2}(\theta^1)^{-1} \right| + \log \Gamma_d\left(\theta^2 + \frac{d+1}{2}\right). \end{aligned} \quad (75)$$

The resulting $D_\alpha^\mathbb{H}(p(x) : q(x))$ and $D_{\alpha, \gamma}^\mathbb{H}(p(x) : q(x))$ are straightforward from the above expression of $F(\theta)$ and Lemma 1. We will omit these tedious expressions for brevity.

4.6 Approximating Hölder projective divergences for statistical mixtures

Given two finite mixture models $m(x) = \sum_{i=1}^k w_i p_i(x)$ and $m'(x) = \sum_{j=1}^{k'} w'_j p'_j(x)$, we derive analytic bounds of their Hölder divergences. When approximation is only needed, one may compute Hölder divergences based on Monte-Carlo stochastic sampling.

Let us assume that all mixture components are in an exponential family [23] so that $p_i(x) = p(x; \theta_i) = \exp(\langle \theta_i, t(x) \rangle - F(\theta_i))$ and $p'_j(x) = p(x; \theta'_j) = \exp(\langle \theta'_j, t(x) \rangle - F(\theta'_j))$ are densities (wrt the Lebesgue measure μ).

We rewrite the Hölder divergence into the form:

$$D_\alpha^\mathbb{H}(m(x) : m'(x)) = -\log \int_{\mathcal{X}} m(x) m'(x) dx + \frac{1}{\alpha} \log \int_{\mathcal{X}} m(x)^\alpha dx + \frac{1}{\beta} \log \int_{\mathcal{X}} m'(x)^\beta dx. \quad (76)$$

To compute the first term, we observe that a product of mixtures is also a mixture:

$$\begin{aligned} \int_{\mathcal{X}} m(x) m'(x) dx &= \sum_{i=1}^k \sum_{j=1}^{k'} k_i k'_j \int_{\mathcal{X}} p_i(x) p'_j(x) dx \\ &= \sum_{i=1}^k \sum_{j=1}^{k'} k_i k'_j \int_{\mathcal{X}} \exp(\langle \theta_i + \theta'_j, t(x) \rangle - F(\theta_i) - F(\theta'_j)) dx \\ &= \sum_{i=1}^k \sum_{j=1}^{k'} k_i k'_j \exp(F(\theta_i) + F(\theta'_j) - F(\theta_i + \theta'_j)), \end{aligned} \quad (77)$$

which can be computed in $O(kk')$ time.

The second and third terms in Eq. 76 are not straightforward to calculate and shall be bounded. Based on computational geometry, we adopt the log-sum-exp bounding technique of [30] and divide the support \mathcal{X} into L pieces of elementary intervals $\mathcal{X} = \bigsqcup_{l=1}^L I_l$. In each interval I_l , the indices:

$$\delta_l = \arg \max_i w_i p_i(x) \text{ and } \epsilon_l = \arg \min_i w_i p_i(x), \quad (78)$$

representing the unique dominating component and the dominated component. Then we bound as follows:

$$\max \left\{ \int_{I_l} k^\alpha w_{\epsilon_l}^\alpha p_{\epsilon_l}(x)^\alpha dx, \int_{I_l} w_{\delta_l}^\alpha p_{\delta_l}(x)^\alpha dx \right\} \leq \int_{I_l} m(x)^\alpha dx \leq \int_{I_l} k^\alpha w_{\delta_l}^\alpha p_{\delta_l}(x)^\alpha dx. \quad (79)$$

All terms on the lhs and rhs of Eq. 79 can be computed exactly by noticing that:

$$\int_I p_i(x)^\alpha dx = \int_I \exp(\langle \alpha \theta_i, t(x) \rangle - \alpha F(\theta_i)) = \exp(F(\alpha \theta_i) - \alpha F(\theta_i)) \int_I p(x; \alpha \theta_i) dx. \quad (80)$$

When $\alpha \theta \in \Theta$ where Θ denotes the natural parameter space, the integral $\int_I p(x; \alpha \theta_i) dx$ converges, see [30] for further details.

Then the bounds of $\int_{\mathcal{X}} m(x)^\alpha dx$ can be obtained by summing the bounds in Eq. 79 over all elementary intervals. Thus $D_{\alpha, \beta}^H(m(x) : m'(x))$ can be both lower and upper bounded.

5 Hölder centroids and center-based clustering

We study the application of HPD and HD for clustering distributions [31], specially clustering Gaussian distributions [35, 12, 24], which have been used in sound processing [35], sensor network [12], statistical debugging [12], and quadratic invariants of switched systems [1], etc. Other potential applications of HD may include nonnegative matrix factorization [37], and clustering von Mises-Fisher [3, 15] (log-normalizer expressed using Bessel functions).

5.1 Hölder centroids

We study center-based clustering of a finite set of distributions: Given a list of distributions belonging to the same conic exponential family with natural parameters $\{\theta_1, \dots, \theta_n\}$ and their associated positive weights $\{w_1, \dots, w_n\}$ with $\sum_{i=1}^n w_i = 1$, consider their centroids based on HPD and HD as follows:

$$C_\alpha(\{\theta_i, w_i\}) = \arg \min_C \sum_{i=1}^n w_i D_\alpha^H(\theta_i : C), \quad (81)$$

$$C_{\alpha, \gamma}(\{\theta_i, w_i\}) = \arg \min_C \sum_{i=1}^n w_i D_{\alpha, \gamma}^H(\theta_i : C). \quad (82)$$

By abuse of notation, C denotes both the HPD centroid and HD centroid. When the context is clear, the parameters in parentheses can be omitted so that these centroids are simply denoted as C_α and $C_{\alpha, \gamma}$. Both of them are defined as the right-sided centroids. The corresponding left-handed centroids are obtained according to reference duality, i.e.,

$$C_{\bar{\alpha}} = \arg \min_C \sum_{i=1}^n w_i D_\alpha^H(C : \theta_i), \quad (83)$$

$$C_{\bar{\alpha}, \gamma} = \arg \min_C \sum_{i=1}^n w_i D_{\alpha, \gamma}^H(C : \theta_i). \quad (84)$$

By Lemma 1, these centroids can be obtained for distributions belonging to the same exponential family as follows:

$$C_\alpha = \arg \min_C \left[\frac{1}{\beta} F(\beta C) - \sum_{i=1}^n w_i F(\theta_i + C) \right], \quad (85)$$

$$C_{\alpha, \gamma} = \arg \min_C \left[\frac{1}{\beta} F(\gamma C) - \sum_{i=1}^n w_i F\left(\frac{\gamma}{\alpha} \theta_i + \frac{\gamma}{\beta} C\right) \right]. \quad (86)$$

Let $\gamma = \alpha$, we get:

$$C_{\alpha, \alpha}(\{\theta_i, w_i\}) = \arg \min_C \left[\frac{1}{\beta} F(\alpha C) - \sum_{i=1}^n w_i F\left(\theta_i + \frac{\alpha}{\beta} C\right) \right] = \frac{\beta}{\alpha} C_\alpha = \frac{1}{\alpha - 1} C_\alpha(\{\theta_i, w_i\}), \quad (87)$$

meaning that the HPD centroid is just a special case of HD centroid up to a *scaling transformation* in the natural parameters space. Let $\gamma = \beta$, we get:

$$C_{\alpha,\beta}(\{\theta_i, w_i\}) = \arg \min_C \left[\frac{1}{\beta} F(\beta C) - \sum_{i=1}^n w_i F\left(\frac{\beta}{\alpha} \theta_i + C\right) \right] = C_\alpha \left(\left\{ \frac{\beta}{\alpha} \theta_i, w_i \right\} \right) = C_\alpha \left(\left\{ \frac{1}{\alpha-1} \theta_i, w_i \right\} \right). \quad (88)$$

Let us consider the general HD centroid $C_{\alpha,\gamma}$. Since F is convex, the minimization energy is the sum of a *convex function* $\frac{1}{\beta} F(\gamma C)$ with a *concave function* $-\sum_{i=1}^n w_i F\left(\frac{\gamma}{\alpha} \theta_i + \frac{\gamma}{\beta} C\right)$. We can therefore use the concave-convex procedure (CCCP) [22] that optimizes Difference of Convex Programs (DCPs): We start with $C_{\alpha,\gamma}^0 = \sum_{i=1}^n w_i \theta_i$ (the barycenter, belonging to Θ) and then update:

$$C_{\alpha,\gamma}^{t+1} = \frac{1}{\gamma} (\nabla F)^{-1} \left(\sum_{i=1}^n w_i \nabla F \left(\frac{\gamma}{\alpha} \theta_i + \frac{\gamma}{\beta} C_{\alpha,\gamma}^t \right) \right). \quad (89)$$

for $t = 0, 1, \dots$ until convergence. This can be done by noting that $\eta = \nabla F(\theta)$ are the dual parameters that are also known as the expectation parameters (or moment parameters). Therefore ∇F and $(\nabla F)^{-1}$ can be computed through Legendre transformations between the natural parameter space and the dual parameter space.

This iterative optimization is guaranteed to converge to a *local* minimum, with a main advantage of bypassing the learning rate parameter of gradient descent algorithms. Since F is strictly convex, ∇F is monotonous, and the rhs expression can be interpreted as a multi-dimensional quasi-arithmetic mean. In fact, it is a barycenter on non-normalized weights scaled by $\beta = \bar{\alpha}$.

For exponential families, the symmetric HPD centroid is:

$$O_\alpha = \arg \min_O \sum_{i=1}^n w_i S_\alpha^H(\theta_i : O) = \arg \min_O \left[\frac{1}{2\alpha} F(\alpha O) + \frac{1}{2\beta} F(\beta O) - \sum_{i=1}^n w_i F(\theta_i + O) \right]. \quad (90)$$

In this case, CCCP update rule is not in closed form because we cannot easily inverse the sum of gradients (but when $\alpha = \beta$, the two terms collapse so CS centroid can be calculated using CCCP). Nevertheless, we can implement the reciprocal operation numerically. Interestingly, the symmetric HD centroid can be solved by CCCP! It amounts to solve:

$$O_{\alpha,\gamma} = \arg \min_O \sum_{i=1}^n w_i S_{\alpha,\gamma}^H(\theta_i : O) = \arg \min_O \left[F(\gamma O) - \sum_{i=1}^n w_i \left(F\left(\frac{\gamma}{\alpha} \theta_i + \frac{\gamma}{\beta} O\right) + F\left(\frac{\gamma}{\beta} \theta_i + \frac{\gamma}{\alpha} O\right) \right) \right]. \quad (91)$$

One can apply CCCP to iteratively update the centroids based on:

$$O_{\alpha,\gamma}^{t+1} = \frac{1}{\gamma} (\nabla F)^{-1} \left[\sum_{i=1}^n w_i \left(\frac{1}{\beta} \nabla F \left(\frac{\gamma}{\alpha} \theta_i + \frac{\gamma}{\beta} O_{\alpha,\gamma}^t \right) + \frac{1}{\alpha} \nabla F \left(\frac{\gamma}{\beta} \theta_i + \frac{\gamma}{\alpha} O_{\alpha,\gamma}^t \right) \right) \right]. \quad (92)$$

Notice the similarity with the updating procedure of $C_{\alpha,\gamma}^t$.

Once the centroid, say $O_{\alpha,\gamma}$, has been computed, we calculate the associated *Hölder information*:

$$\sum_{i=1}^n w_i S_{\alpha,\gamma}^H(\theta_i : O_{\alpha,\gamma}). \quad (93)$$

The Hölder information generalizes the notion of variance and Bregman information [4] to the case of Hölder distances.

5.2 Clustering based on symmetric Hölder divergences

Given a set of fixed densities $\{p_1, \dots, p_n\}$, we can perform *variational k-means* [28] with respect to the Hölder divergence to minimize the cost function:

$$E(O_1, \dots, O_L, l_1, \dots, l_n) = \sum_{i=1}^n S_{\alpha,\gamma}^H(p_i : O_{l_i}), \quad (94)$$

where O_1, \dots, O_L are the cluster centers, and $l_i \in \{1, \dots, L\}$ is the cluster label of p_i . The algorithm is given by Algorithm 1. Notice that one does not need to wait for the CCCP iterations to converge. It only has to improve the cost function E before updating the assignment. We have implemented the algorithm based on the symmetric HD. One can easily modify it based on HPD and other variants.

Algorithm 1: Hölder variational k -means

Input: A list of probability distributions p_1, \dots, p_n ; number of clusters L ; $\alpha > 1$; $\gamma > 0$
Output: A clustering scheme $p_i \rightarrow \{1, \dots, L\}$, $\forall i \in \{1, \dots, n\}$

```

1 Randomly assign  $p_i$  to  $l_i$ 
2 while not converged do
3   for  $l = 1, \dots, L$  do
4     /* Variational  $k$ -means: Carry CCCP iterations until the current center
       improves the former cluster Hölder information */
     Compute the centroid  $O_l = \arg \min_O \sum_{i:l_i=l} S_{\alpha,\gamma}^H(p_i : O)$ 
5   for  $i = 1, \dots, n$  do
6     Assign  $l_i = \arg \min_l S_{\alpha,\gamma}^H(p_i : O_l)$ 
7 return  $\{l_i\}_{i=1}^n$ 

```

Table 1: Clustering accuracy of the 2D Gaussian dataset (based on 500 independent runs).


Parameters	$\alpha = \gamma = 1.1$	$\alpha = \gamma = 1.5$	$\alpha = \gamma = 2$ (CS)	$\alpha = \gamma = 10$
$n = 50$	95.6% \pm 7.6%	92.4 \pm 9.8%	92.2% \pm 9.9%	92.2% \pm 9.7%
$n = 100$	97.3% \pm 5.3%	94.9 \pm 8.8%	94.0% \pm 10.1%	94.2% \pm 9.3%

We made a toy dataset generator which can randomly generate n 2D Gaussians which have a underlying structure of two clusters. In the first cluster, the mean of each Gaussian $G(\mu, \Sigma)$ has the prior distribution $\mu \sim G((-2, 0), I)$; the covariance matrix is obtained by first generating $\sigma_1 \sim \Gamma(7, 0.01)$, $\sigma_2 \sim \Gamma(7, 0.003)$, where Γ means a gamma distribution with prescribed shape and scale, then rotating the covariance matrix $\text{diag}(\sigma_1, \sigma_2)$ so that the resulting Gaussian has a “radial direction” with respect to the center $(-2, 0)$. The second cluster is similar to the first cluster with the only difference being that $\mu \sim G((2, 0), I)$ is centered around $(2, 0)$. See Fig. 4 for an intuitive presentation of the toy dataset.


To reduce the number of parameters that has to be tuned, we only investigate the case $\alpha = \gamma$. If we choose $\alpha = \gamma = 2$, then $S_{\alpha,\gamma}^H$ becomes the CS divergence and Algorithm 1 reduces to traditional CS clustering. From Fig. 4 we can observe that the clustering result does vary with the settings of α and γ . We performed clustering experiments on two different settings of the sample size. Table 1 shows the clustering accuracy measured by the percentage of “correctly clustered” Gaussians, i.e., the output label by clustering algorithms that coincides with the true label corresponding to the data generating process. We see that the symmetric Hölder divergence can give better clustering results as compared to CS clustering. This hints that one could consider the general Hölder divergence to replace CS in similar clustering applications [16, 33]. Although one faces the problem of tuning the parameter α and γ , Hölder divergences can potentially give *better* results. This is expected because CS is just one particular case of the class of Hölder divergences.

6 Conclusion and perspectives

We introduced the notion of pseudo-divergences that generalizes the concept of divergences in information geometry [2] that are smooth non-metric statistical distances that are not required to obey the law of the indiscernibles. Pseudo-divergences can be built from inequalities by considering the inequality difference gap or its log-ratio gap. We then defined two classes of statistical measures based on Hölder’s ordinary and reverse inequalities: The singly parametric family of Hölder pseudo-divergences and the doubly parametric family of Hölder divergences. By construction, the Hölder divergences are proper divergences between probability densities. Both statistical Hölder distance families are projective divergences that



(a) $\alpha = \gamma = 1.1$ (Hölder clustering)



(b) $\alpha = \gamma = 2$ (Cauchy-Schwarz clustering)

Figure 4: Variational k -means clustering results on a toy dataset consisting of a set of 2D Gaussians organized into two clusters.

do not require distributions to be normalized, and admit closed-form expressions when considering exponential families with conic or affine natural parameter space (like multinomials or multivariate normals). Those two families of distances can be symmetrized and intersect at the unique Cauchy-Schwarz divergence. Since the Cauchy-Schwarz divergence is often used in distribution clustering applications [16], we carried out preliminary experiments demonstrating experimentally that the symmetrized Hölder divergences improved over the Cauchy-Schwarz divergence for a toy dataset of Gaussians. We briefly touched upon the use of these novel divergences in statistical estimation theory. These projective Hölder (pseudo-)divergences are different from the recently introduced composite scored-induced Hölder di-

vergences [20, 19] that are not projective divergences and do not admit closed-form expressions for exponential families in general.

We elicited the special role of escort distributions [2] for Hölder divergences in our framework: Escort distributions transform distributions to allow one:

- To reveal that Hölder pseudo-divergences on escort distributions amount to skew Bhattacharyya divergences [22],
- To transform the improper Hölder pseudo-divergences into proper Hölder divergences, and vice versa.

Let us conclude with a perspective note on pseudo-divergences, statistical estimators, and manifold learning. Proper divergences have been widely used in statistical estimators to build families of estimators [32, 5]. Similarly, given a prescribed density $p_0(x)$, a pseudo-divergence yields a corresponding estimator by minimizing $D(p_0(x) : q(x))$ with respect to $q(x)$. However in this case the resulting $q(x)$ is potentially biased and is not guaranteed to recover the optimal input $p_0(x)$. Furthermore, the minimizer of $D(p_0(x) : q(x))$ may not be unique, i.e., there could be more than one probability density $q(x)$ yielding $D(p_0(x) : q(x)) = 0$.

How can pseudo-divergences be useful? We have the following two simple arguments:

- In an estimation scenario, we can usually pre-compute $p_1(x) \neq p_0(x)$ according to $D(p_1(x) : p_0(x)) = 0$. Then the estimation $q(x) = \arg \min_q D(p_1(x) : q(x))$ will automatically target at $p_0(x)$. We call this technique *pre-aim*.

For example, given positive measure $p(x)$, we first find $p_0(x)$ to satisfy $D_\alpha^H(p_0(x) : p(x)) = 0$. We have $p_0(x) = p(x)^{\frac{1}{\alpha-1}}$ that satisfies this condition. Then a proper divergence between $p(x)$ and $q(x)$ can be obtained by aiming $q(x)$ towards $p_0(x)$. For conjugate exponents α and β ,

$$D_\alpha^H(p_0(x) : q(x)) = -\log \frac{\int_{\mathcal{X}} p_0(x) q(x) dx}{\left(\int_{\mathcal{X}} p_0(x)^\alpha dx\right)^{1/\alpha} \left(\int_{\mathcal{X}} q(x)^\beta dx\right)^{1/\beta}}, \quad (95)$$

$$= -\log \frac{\int_{\mathcal{X}} p(x)^{\frac{1}{\alpha-1}} q(x) dx}{\left(\int_{\mathcal{X}} p(x)^{\frac{\alpha}{\alpha-1}} dx\right)^{1/\alpha} \left(\int_{\mathcal{X}} q(x)^\beta dx\right)^{1/\beta}}, \quad (96)$$

$$= -\log \frac{\int_{\mathcal{X}} p(x)^{\frac{\beta}{\alpha}} q(x) dx}{\left(\int_{\mathcal{X}} p(x)^\beta dx\right)^{1/\alpha} \left(\int_{\mathcal{X}} q(x)^\beta dx\right)^{1/\beta}} = D_{\alpha,\beta}^H(p(x) : q(x)). \quad (97)$$

This means that the pre-aim technique of HPD is equivalent to HD $D_{\alpha,\gamma}^H$ when we set $\gamma = \beta$.

As an alternative implementation of pre-aim, since $D_\alpha^H(p(x) : p^{\alpha-1}(x)) = 0$, a proper divergence between $p(x)$ and $q(x)$ can be constructed by measuring:

$$D_\alpha^H(q(x) : p^{\alpha-1}(x)) = -\log \frac{\int_{\mathcal{X}} p(x)^{\frac{\alpha}{\beta}} q(x) dx}{\left(\int_{\mathcal{X}} q(x)^\alpha dx\right)^{1/\alpha} \left(\int_{\mathcal{X}} p(x)^\alpha dx\right)^{1/\beta}} = D_{\beta,\alpha}^H(p(x) : q(x)), \quad (98)$$

turning out again to belong to the class of HD.

In practice, HD as a two-parameter family is less used than HPD with pre-aim because of the difficulty in choosing the parameter γ , and because that HD has a slightly more complicated expression. The family of HD connecting CS divergence with skew Bhattacharyya divergence [22] is nevertheless of theoretical importance.

- In manifold learning [17, 10, 38], it is an essential topic to align two category distributions $p_0(x)$ and $q(x)$ corresponding respectively to the input and output [38], both for learning and for performance evaluation. In this case, the dimensionality of the statistical manifold that encompasses $p_0(x)$ and $q(x)$ is so high that to preserve monotonically $p_0(x)$ in the resulting $q(x)$ is already a difficult non-linear optimization and could be sufficient for the application, while preserving perfectly the input $p_0(x)$ is not so meaningful because of the input noise. It is then much easier to define pseudo-divergences using inequalities which are not necessarily to be proper with potentially more choices.

On the other hand, projective divergences including Hölder divergences introduced in this work are more meaningful in manifold learning than KL divergence (which is widely used) because they give scale invariance of the probability densities, meaning that one can define positive similarities then directly align these similarities, which is guaranteed to be equivalent to align the corresponding distributions. This could potentially give unified perspectives in between the two approaches of similarity based manifold learning [10] and the probabilistic approach [17].

We expect that these two novel parametric Hölder classes of statistical divergences and pseudo-divergences open up new insights and applications in statistics and information sciences. Furthermore, the framework to build divergences or pseudo-divergences from proper or improper biparametric inequalities offers novel classes of divergences to study.

Reproducible source code is available online at:
<https://www.lix.polytechnique.fr/~nielsen/HPD/>

References

- [1] Xavier Allamigeon, Stéphane Gaubert, Eric Goubault, Sylvie Putot, and Nikolas Stott. A scalable algebraic method to infer quadratic invariants of switched systems. In *Proceedings of the 12th International Conference on Embedded Software*, pages 75–84. IEEE Press, 2015.
- [2] Shun-ichi Amari. *Information Geometry and Its Applications*. Applied Mathematical Sciences. Springer Japan, 2016.
- [3] Arindam Banerjee, Inderjit S Dhillon, Joydeep Ghosh, and Suvrit Sra. Clustering on the unit hypersphere using von Mises-Fisher distributions. *Journal of Machine Learning Research*, 6(Sep):1345–1382, 2005.
- [4] Arindam Banerjee, Srujana Merugu, Inderjit S Dhillon, and Joydeep Ghosh. Clustering with Bregman divergences. *Journal of machine learning research*, 6(Oct):1705–1749, 2005.
- [5] Ayanendranath Basu, Hiroyuki Shioya, and Chanseok Park. *Statistical inference: the minimum distance approach*. CRC Press, 2011.
- [6] A. Bhattacharyya. On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, 35:99–109, 1943.
- [7] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [8] Marcin Budka, Bogdan Gabrys, and Katarzyna Musiał. On accuracy of PDF divergence estimators and their applicability to representative data sampling. *Entropy*, 13(7):1229–1266, 2011.
- [9] J Burbea and C Rao. On the convexity of some divergence measures based on entropy functions. *IEEE Transactions on Information Theory*, 28(3):489–495, 1982.
- [10] M. Á. Carreira-Perpiñán. The elastic embedding algorithm for dimensionality reduction. In *International Conference on Machine Learning*, pages 167–174, 2010.
- [11] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [12] Jason V Davis and Inderjit S Dhillon. Differential entropic clustering of multivariate Gaussians. In *Advances in Neural Information Processing Systems*, pages 337–344, 2006.
- [13] Hironori Fujisawa and Shinto Eguchi. Robust parameter estimation with a small bias against heavy contamination. *Journal of Multivariate Analysis*, 99(9):2053–2081, 2008.
- [14] Tilman Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [15] Siddarth Gopal and Yiming Yang. Von Mises-Fisher clustering models. *Journal of Machine Learning Research*, 32:154–162, 2014.

- [16] Erion Hasanbelliu, Luis Sanchez Giraldo, and Jose C Principe. Information theoretic shape matching. *IEEE transactions on pattern analysis and machine intelligence*, 36(12):2436–2451, 2014.
- [17] G. E. Hinton and S. T. Roweis. Stochastic Neighbor Embedding. In *Advances in Neural Information Processing Systems 15 (NIPS 2002)*, pages 833–840. MIT Press, 2003.
- [18] Otto Ludwig Hölder. Über einen Mittelwertssatz. *Nachr. Akad. Wiss. Gottingen Math.-Phys. Kl.*, 1889.
- [19] Takafumi Kanamori. Scale-invariant divergences for density functions. *Entropy*, 16(5):2611–2628, 2014.
- [20] Takafumi Kanamori, Hironori Fujisawa, et al. Affine invariant divergences associated with proper composite scoring rules and their applications. *Bernoulli*, 20(4):2278–2304, 2014.
- [21] Frank Nielsen. Closed-form information-theoretic divergences for statistical mixtures. In *Pattern Recognition (ICPR), 2012 21st International Conference on*, pages 1723–1726. IEEE, 2012.
- [22] Frank Nielsen and Sylvain Boltz. The Burbea-Rao and Bhattacharyya centroids. *IEEE Transactions on Information Theory*, 57(8):5455–5466, 2011.
- [23] Frank Nielsen and Vincent Garcia. Statistical exponential families: A digest with flash cards. *arXiv preprint arXiv:0911.4863*, 2009.
- [24] Frank Nielsen and Richard Nock. Clustering multivariate normal distributions. In *Emerging Trends in Visual Computing*, pages 164–174. Springer, 2009.
- [25] Frank Nielsen and Richard Nock. A closed-form expression for the Sharma-Mittal entropy of exponential families. *Journal of Physics A: Mathematical and Theoretical*, 45(3):032003, 2011.
- [26] Frank Nielsen and Richard Nock. Skew Jensen-Bregman Voronoi diagrams. In *Transactions on Computational Science XIV*, pages 102–128. Springer, 2011.
- [27] Frank Nielsen and Richard Nock. On the Chi square and higher-order Chi distances for approximating f -divergences. *IEEE Signal Processing Letters*, 1(21):10–13, 2014.
- [28] Frank Nielsen and Richard Nock. Total Jensen divergences: definition, properties and clustering. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2016–2020. IEEE, 2015.
- [29] Frank Nielsen and Richard Nock. Patch matching with polynomial exponential families and projective divergences. In *Similarity Search and Applications - 9th International Conference, SISAP 2016, Tokyo, Japan, October 24-26, 2016. Proceedings*, pages 109–116, 2016.
- [30] Frank Nielsen and Ke Sun. Guaranteed bounds on information-theoretic measures of univariate mixtures using piecewise log-sum-exp inequalities. *Entropy*, 18(12):442, 2016.
- [31] Akifumi Notsu, Osamu Komori, and Shinto Eguchi. Spontaneous clustering via minimum gamma-divergence. *Neural computation*, 26(2):421–448, 2014.
- [32] Leandro Pardo. *Statistical Inference Based on Divergence Measures*. CRC Press, Abingdon, 2005.
- [33] Hassan Rami, Leila Belmerhnia, Ahmed Drissi El Maliani, and Mohammed El Hassouni. Texture retrieval using mixtures of generalized Gaussian distribution and Cauchy-Schwarz divergence in wavelet domain. *Image Commun.*, 42(C):45–58, March 2016.
- [34] C. R. Rao. Information and accuracy attainable in the estimation of statistical parameters. *Bulletin of Calcutta Mathematical Society*, 37(3):81–91, 1945.
- [35] Luca Rigazio, Brice Tsakam, and J-C Junqua. An optimal Bhattacharyya centroid algorithm for Gaussian clustering with applications in automatic speech recognition. In *Acoustics, Speech, and Signal Processing, 2000. ICASSP'00. Proceedings. 2000 IEEE International Conference on*, volume 3, pages 1599–1602. IEEE, 2000.

- [36] Anuj Srivastava, Ian Jermyn, and Shantanu Joshi. Riemannian analysis of probability density functions with applications in vision. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007.
- [37] Dennis L Sun and Cédric Févotte. Alternating direction method of multipliers for non-negative matrix factorization with the beta-divergence. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6201–6205. IEEE, 2014.
- [38] Ke Sun and Stéphane Marchand-Maillet. An information geometry of statistical manifold learning. In *International Conference on Machine Learning; JMLR W&CP 32 (2)*, pages 1–9, 2014.
- [39] Jun Zhang. Reference duality and representation duality in information geometry. In *American Institute of Physics Conference Series*, volume 1641, pages 130–146, 2015.

A Proof of Hölder ordinary and reverse inequalities

We extend the proof ([7], pp.78) to prove both the (ordinary or forward) Hölder inequality and the reverse Hölder inequality.

Proof. First, let us observe that $-\log(x)$ is strictly convex on $(0, +\infty)$ since $(-\log(x))'' = \frac{1}{x^2}$. It follows that for $0 < a < 1$ that:

$$-\log(ax_1 + (1-a)x_2) \leq -a\log(x_1) - (1-a)\log(x_2), \quad (99)$$

where the equality holds iff $x_1 = x_2$.

Conversely, when $a < 0$ or $a > 1$, we have:

$$-\log(ax_1 + (1-a)x_2) \geq -a\log(x_1) - (1-a)\log(x_2), \quad (100)$$

where the equality holds iff $x_1 = x_2$.

Equivalently, we can write these two inequalities as follows:

$$\begin{cases} x_1^a x_2^{1-a} \leq ax_1 + (1-a)x_2 & \text{(if } 0 < a < 1\text{);} \\ x_1^a x_2^{1-a} \geq ax_1 + (1-a)x_2 & \text{(if } a < 0 \text{ or } a > 1\text{),} \end{cases} \quad (101)$$

both of them are tight iff $x_1 = x_2$.

Let P and Q be positive measures with Radon-Nikodym densities $p(x) > 0$ and $q(x) > 0$ be positive densities with respect to the reference Lebesgue measure μ . The densities are strictly greater than 0 on the support \mathcal{X} . Plugging:

$$a = \frac{1}{\alpha}, \quad 1-a = \frac{1}{\beta}, \quad x_1 = \frac{p(x)^\alpha}{\int_{\mathcal{X}} p(x)^\alpha dx}, \quad x_2 = \frac{q(x)^\beta}{\int_{\mathcal{X}} q(x)^\beta dx}, \quad (102)$$

into Eq. 101, we get:

$$\begin{cases} \frac{p(x)}{(\int_{\mathcal{X}} p(x)^\alpha dx)^{1/\alpha}} \frac{q(x)}{(\int_{\mathcal{X}} q(x)^\beta dx)^{1/\beta}} \leq \frac{1}{\alpha} \frac{p(x)^\alpha}{\int_{\mathcal{X}} p(x)^\alpha dx} + \frac{1}{\beta} \frac{q(x)^\beta}{\int_{\mathcal{X}} q(x)^\beta dx} & \text{if } \alpha > 0 \text{ and } \beta > 0, \\ \frac{p(x)}{(\int_{\mathcal{X}} p(x)^\alpha dx)^{1/\alpha}} \frac{q(x)}{(\int_{\mathcal{X}} q(x)^\beta dx)^{1/\beta}} \geq \frac{1}{\alpha} \frac{p(x)^\alpha}{\int_{\mathcal{X}} p(x)^\alpha dx} + \frac{1}{\beta} \frac{q(x)^\beta}{\int_{\mathcal{X}} q(x)^\beta dx} & \text{if } \alpha < 0 \text{ or } \beta < 0. \end{cases} \quad (103)$$

Assume that $p(x)$ in $L_\alpha(\mathcal{X}, \mu)$ and $q(x)$ in $L_\beta(\mathcal{X}, \mu)$ so that both $\int_{\mathcal{X}} p(x)^\alpha dx$ and $\int_{\mathcal{X}} q(x)^\beta dx$ converge. Integrate both sides on \mathcal{X} to get:

$$\begin{cases} \frac{\int_{\mathcal{X}} p(x)q(x)dx}{(\int_{\mathcal{X}} p(x)^\alpha dx)^{1/\alpha} (\int_{\mathcal{X}} q(x)^\beta dx)^{1/\beta}} \leq 1 & \text{if } \alpha > 0 \text{ and } \beta > 0, \\ \frac{\int_{\mathcal{X}} p(x)q(x)dx}{(\int_{\mathcal{X}} p(x)^\alpha dx)^{1/\alpha} (\int_{\mathcal{X}} q(x)^\beta dx)^{1/\beta}} \geq 1 & \text{if } \alpha < 0 \text{ or } \beta < 0. \end{cases} \quad (104)$$

The necessary and sufficient condition for equality is that:

$$\frac{p(x)^\alpha}{\int_{\mathcal{X}} p(x)^\alpha dx} = \frac{q(x)^\beta}{\int_{\mathcal{X}} q(x)^\beta dx}, \quad (105)$$

almost everywhere. That is, there exists a positive constant $\lambda > 0$ such that:

$$p(x)^\alpha = \lambda q(x)^\beta, \quad \lambda > 0, \quad \text{almost everywhere.} \quad (106)$$

□

The Hölder conjugate exponents α and β satisfies $\frac{1}{\alpha} + \frac{1}{\beta} = 1$. That is $\beta = \frac{\alpha}{\alpha-1}$. Thus when $\alpha < 0$ we necessarily have $\beta > 0$, and vice-versa.

We can unify these two straight and reverse Hölder inequalities into a single inequality by considering the sign of $\alpha\beta = \frac{\alpha^2}{\alpha-1}$: We get the general Hölder inequality:

$$(-1)^{\text{sign}(\alpha\beta)} \frac{\int_{\mathcal{X}} p(x)q(x)dx}{\left(\int_{\mathcal{X}} p(x)^\alpha dx\right)^{1/\alpha} \left(\int_{\mathcal{X}} q(x)^\beta dx\right)^{1/\beta}} \geq (-1)^{\text{sign}(\alpha\beta)}. \quad (107)$$

When $\alpha = \beta = 2$, Hölder inequality becomes the Cauchy-Schwarz inequality:²

$$\int_{\mathcal{X}} p(x)q(x)dx \leq \sqrt{\left(\int_{\mathcal{X}} p(x)^2 dx\right) \left(\int_{\mathcal{X}} q(x)^2 dx\right)}. \quad (108)$$

B Notations

D_{α}^{HS}	Hölder proper non-projective Scored-induced divergence [19]
D_{α}^{H}	Hölder improper projective pseudo-divergence (new)
$D_{\alpha,\gamma}^{\text{H}}$	Hölder proper projective divergence (new)
D_{α}^{HE}	Hölder proper projective escort divergence (new)
KL	Kullback-Leibler divergence [11]
CS	Cauchy-Schwarz divergence [8]
B	Bhattacharyya distance [6]
$B_{\frac{1}{\alpha}}$	skew Bhattacharyya distance [22]
D_{γ}	γ -divergence (score-induced) [13]
$p_{\alpha}^E, q_{\beta}^E$	escort distributions
α, β	Hölder conjugate pair of exponents: $\frac{1}{\alpha} + \frac{1}{\beta} = 1$
$\bar{\alpha}, \bar{\beta}$	Hölder conjugate exponent: $\bar{\alpha} = \beta = \frac{\alpha}{\alpha-1}$
θ_p, θ_q	natural parameters of exponential family distributions
\mathcal{X}	support of distributions
μ	Lebesgue measure
$L^{\gamma}(\mathcal{X}, \mu)$	Lebesgue space of functions f such that $\int_{\mathcal{X}} f(x) ^{\gamma} dx < \infty$

Contents

1	Introduction: Inequality, proper divergence and improper pseudo-divergence	1
1.1	Statistical divergences from inequality gaps	1
1.2	Pseudo-divergences and the axiom of indiscernability	3
1.3	Prior work and contributions	3
1.4	Organization	4
2	Hölder pseudo-divergence: Definition and properties	4
2.1	Definition	4
2.2	Properness and improperness	5
2.3	Reference duality	6
2.4	HPD is a projective divergence	6
2.5	Escort distributions and skew Bhattacharyya divergences	6

²Historically, Cauchy stated the discrete sum inequality in 1821 while Schwarz reported the integral form of the inequality in 1888.

3	Proper Hölder divergence	7
3.1	Definition	7
3.2	Special case: The Cauchy-Schwarz divergence	7
3.3	Limit cases of Hölder divergences and statistical estimation	8
4	Closed-form expressions of HPD and HD for conic and affine exponential families	9
4.1	Case study: Categorical distributions	11
4.2	Case study: Bernoulli distribution	13
4.3	Case study: MultiVariate Normal distributions (MVNs)	13
4.4	Case study: Zero-centered Laplace distribution	14
4.5	Case study: Wishart distribution	15
4.6	Approximating Hölder projective divergences for statistical mixtures	15
5	Hölder centroids and center-based clustering	16
5.1	Hölder centroids	16
5.2	Clustering based on symmetric Hölder divergences	17
6	Conclusion and perspectives	18
A	Proof of Hölder ordinary and reverse inequalities	23
B	Notations	24